AD-A074 956 STANFORD UNIV CALIF DEPT OF ENGINEERING-ECONOMIC SYSTEMS F/G 12/1
OPTIMAL INFORMATION ACQUISITION FOR RANDOMLY OCCURRING DECISION--ETC(U)
SEP 79 A SHARIFNIA N00014-79-C-0036
UNCLASSIFIED EES-DA-79-1 NL

1 OF 2
AD A 074956

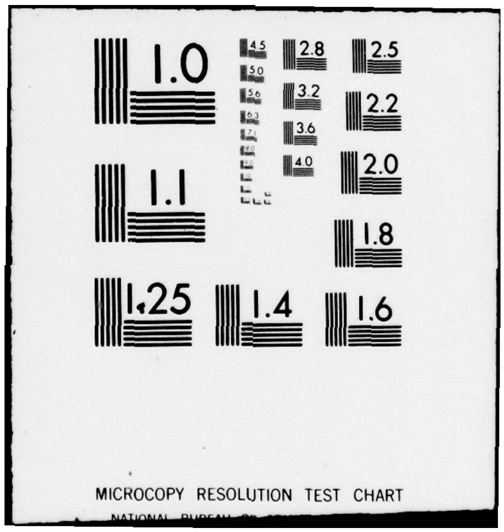
1.0
1.1
1.25
1.4
1.6
1.8
2.0
2.2
2.5
2.8
3.2
3.6
4.0
4.5
5.0
5.6
6.3
MICROCOPY RESOLUTION TEST CHART

Research Report No. EES-DA-79-1
September 1979

OPTIMAL INFORMATION ACQUISITION FOR RANDOMLY OCCURRING DECISIONS

ALIMOHAMMAD SHARIFNIA

DECISION ANALYSIS PROGRAM

Professor Ronald A. Howard
Principal Investigator

DEPARTMENT OF ENGINEERING-ECONOMIC SYSTEMS

Stanford University
Stanford, California 94305

SPONSORSHIP

Advanced Research Projects Agency, Cybernetics Technology Office, monitored by Engineering Psychology Programs, Office of Naval Research, Contract No. N00014-79-C-0036.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER (14) EES-DA-79-1	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) (6) Optimal Information Acquisition for Randomly Occurring Decisions.		5. TYPE OF REPORT & PERIOD COVERED (9) Technical rept.,
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) (10) Alimohammad Sharifnia		8. CONTRACT OR GRANT NUMBER(s) (15) N00014-79-C-0036 new
9. PERFORMING ORGANIZATION NAME AND ADDRESS The Board of Trustees of the Leland Stanford Junior University, c/o Office of Research Administrator, Encina Hall, Stanford, Ca. 94305		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS (12) 155
11. CONTROLLING OFFICE NAME AND ADDRESS Advanced Research Projects Agency Cybernetics Technology Office 1400 Wilson Blvd., Arlington, Va. 22209		12. REPORT DATE (11) Sep[redacted] 1979
		13. NUMBER OF PAGES 154
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Engineering Psychology Programs Office of Naval Research 800 No. Quincy St., Arlington, Va. 22217		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release. Distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

INFORMATION PERISHING — DYNAMIC DECISIONS

INFORMATION RECOVERY

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This research investigates the information acquisition policies for randomly occurring decisions. A randomly occurring decision is a decision that must be made upon the occurrence of a precipitating event that occurs randomly with time. Due to the urgency often associated with this type of decision, it is difficult to obtain fresh information at the time of the decision. Therefore, the available over

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE S/N 0102-LF-014-6601

UNCLASSIFIED 408 566 LB

information will be limited to the unfresh (old) information from the past. Since the information becomes outdated and obsolete in time, regular updating of the information is often desirable in order to be prepared for the decision.

The process of information outdating (perishing) and its relationship with the characteristics of the dynamic environment, the decision for which the information is used, and the type of information (perfect or imperfect) are investigated. Based on this process, as well as the cost of information recovery (updating), and the likelihood of the occurrence of the decision, the optimal policies for the recovery of information are studied. Policies which use only the prior knowledge about the environment as well as those which utilize the information in each observation (in addition to the prior knowledge) are analyzed, and optimality conditions are found for each case. The case of a one-time decision, namely when the decision happens only once, is studied initially. The results are then extended to the case where the decision may be repeated in time.

ABSTRACT

This research investigates the information acquisition policies for randomly occurring decisions. A randomly occurring decision is a decision that must be made upon the occurrence of a precipitating event that occurs randomly with time. Due to the urgency often associated with this type of decision it is difficult to obtain fresh information at the time of the decision. Therefore, the available information will be limited to the unfresh (old) information from the past. Since the information becomes outdated and obsolete in time, regular updating of the information is often desirable in order to be prepared for the decision.

The process of information outdating (perishing) and its relationship with the characteristics of the dynamic environment, the decision for which the information is used, and the type of information (perfect or imperfect) are investigated. Based on this process, as well as the cost of information recovery (updating), and the likelihood of the occurrence of the decision, the optimal policies for the recovery of information are studied. Policies which use only the prior knowledge about the environment as well as those which utilize the information in each observation (in addition to the prior knowledge) are analyzed, and optimality conditions are found for each case. The case of a one-time decision, namely when the decision happens only once, is studied initially. The results are then extended to the case where the decision may be repeated in time.

TABLE OF CONTENTS

TABLE OF CONTENTS (Cont'd)

TABLE OF CONTENTS (Cont'd)

LIST OF TABLES

LIST OF ILLUSTRATIONS

LIST OF ILLUSTRATIONS (Cont'd)

LIST OF ILLUSTRATIONS (Cont'd)

CHAPTER 1

INTRODUCTION

1.1 Introduction and Background

The role of information in the modern society is of the utmost importance. The existence of large social and economic organizations which characterize today's advanced societies depends on the capabilities for producing, processing, updating, and communicating information on a very large scale. It is well established that in an advanced society a great portion of the labor force is involved in activities related to information, and a high percentage of the cost of running the economy consists of information related costs [13].* Yet the quantitative analysis to rationalize the allocation of resources to the various activities related to the information process has been very limited.

The most important problem in dealing with information is,perhaps, the problem of measurement. We often have an intuitive tendency to regard information as a commodity and to apply the notions of "more" and "less" to it just as we apply them to commodities. Upon more careful examination, however, we find that there is no clear way we can attribute such notions to information in a general and precise fashion [14]. Therefore, we need an appropriate context in which information can be treated in a meaningful, proper, and precise manner. Decision analysis provides such a context. Decision analysis treats information explicitly and evaluates it on the basis of its economic value for making decisions. Various pieces of information are compared according to their economic values for making a specific decision. It is only natural, therefore, to find that one

---

*Numbers in square brackets refer to the References found at the end of the thesis.

piece of information is more valuable than another for one decision, while the opposite is true for a different decision. This would be inconsistent if measures of information (e.g. "quantity of information") are defined in the abstract, separate from the decision itself.

Previous studies of information in the context of decision analysis have been mostly static, in the sense that they make no reference to the time of the decision or the age of the information. In many real-world cases, however, information has some kind of time association. Consider, for example, a large organization in which information related to the activities of the organization is accumulated and updated for making future decisions. Here we have a continuing process of information production and updating for the general task of decision making in the organization. In such a scenario, information is dynamic and must be studied in a dynamic framework.

An important case in which a dynamic framework is essential for the study of information is when the time of the decision is uncertain. The decision must be made upon the occurrence of a precipitating event which occurs randomly with time. The decision maker often has no control over this event. This type of decision is referred to as a contingent decision [5]. We can find numerous examples of contingent decisions: firms have to make contingent decisions in response to action by their competitors or by the government (at uncertain times); the government often faces contingent decisions as a result of other countries' economic or political decisions.

Due to the urgency often associated with this type of decision, there is typically insufficient time to obtain adequate information for

the decision after the precipitating event has occurred. It is therefore desirable to be prepared for the decision by obtaining the necessary information in advance. A problem arises, however, because we are generally faced with a changing environment in which the acquired information becomes outdated and obsolete as time passes. Suppose, for example, that we are expecting a contingent decision and that the level of an inventory (at the decision time) is a valuable piece of information for making the decision. If we learn the level of the inventory today, this information would be very valuable if the decision happens today, but it may not be as valuable if the decision happens a week from now. We may expect that this information would be less and less valuable as time passes. This is called information outdating or perishing. Faced with this problem, we need to update our information regularly in order to remain prepared for the decision. This is referred to as information recovery or replenishment. In the above example regular observations of the level of the inventory would enhance our information at the time of the decision. The more frequently the observations are made, the more accurate our information will be at the decision time. But the information is costly and therefore we have to balance the costs and the benefits of information in order to find the optimum policies for the recovery of information.

It is clear from the above discussion that the important problems with regard to information in a dynamic framework are the outdating and the recovery of information. These problems are analyzed in this study.

There are other decisions which are in many respects similar to contingent decisions. As in the case of a contingent decision, there

is uncertainty about the time of the decision and therefore information must be obtained in advance. But there is also uncertainty in the exact nature of the decision. Consider, for example, all the data gathered and updated by the government or private agencies in anticipation of future use. There is uncertainty regarding the time of use as well as the purpose (decision) for which the data will be used (although the general area of use may be known). We will confine our study to contingent decisions, namely the case in which there is no ambiguity in the decision itself. The results may be useful, however, for the study of information acquisition policies for the second type of problem.

There has been surprisingly little research in this area. A preliminary work by Marschak and Radner [11] gives a valuable formulation of the problem but does not establish concrete results on the dyanmics of information. This work was motivated by a recent dissertation by Grum [5] who studied the "perishing" and "replenishment" of information in a dynamic environment. His analysis is, however, limited to discrete-state dynamic systems (Markov chains). In this research the problem is formulated in the framework of continuous-state systems (following Marschak and Radner), which allows a more systematic study of the problem and facilitates the derivation of more general results.

1.2 A Contingent-Decision Example

This example illustrates the problems of information outdating and recovery for a contingent decision, and it is studied in detail in subsequent chapters.

Suppose that we are expecting a bidding occasion but are uncertain of its time. When the bidding is announced, we will have to make our bid within a short time and, therefore, will not be able to obtain new information on the uncertain variables at the bidding time. An important variable for the bidding is our cost of performing the contract (p), and clearly, information on this variable would help us make a better bid. We can find p (at some cost) at any time. However, since this variable changes over time, knowledge of its present value will not provide us with perfect information about its value at a future time, in particular at the time of the bidding. We expect that the older our information at the bidding time, the less valuable it will be (information perishing). Note that we are implicitly assuming a relationship among the values of p at different points in time, and that we have some knowledge of this relationship from past experience. If this was not the case, current knowledge of p would have no value for making the bid in the future.

Suppose, for example, that from past experience we know that P has a constant mean over time, and that its variation from its mean (ΔP) changes over time according to the linear Markovian model

$$\Delta p(t) = \lambda \cdot \Delta p(t-1) + \varepsilon(t) \qquad (1)$$

where λ is a constant and $\varepsilon(t)$ is a random "noise" term with zero mean. This knowledge makes it possible to use our present information about p when making our bid at a future time.

Suppose that we have obtained information about the cost of performing the contract at present, $(p(0))$. As time passes, this

information becomes old and we must decide when to update it. The optimum updating time depends on how quickly the existing information perishes, on the cost of a new observation of p, and on the likelihood of the bidding occurring at each future point in time. The perishing of information depends not only on the dynamics of $p(t)$, (Eq. (1)), but also on the decision for which the information will be used (bidding payoff), and on the type of information (perfect or imperfect). The optimum information recovery time may also depend on the result of our previous observation of $p(t)$, namely $p(0)$. These are the types of problems which are investigated in this research. The analysis is generalized as far as possible in order to facilitate the application of the results to various specific cases.

1.3 Overview

In Chapter 2 the dynamics of information in a dynamic environment are investigated. It is shown how the dynamics of information are influenced by the factors which influence them, namely the characteristics of the dynamic environment, the decision for which the information is used, and the type of information itself (perfect or imperfect). The possible patterns of value of information in time under various conditions are investigated.

In Chapter 3 we investigate the problem of information recovery (updating) in the face of a dynamic environment, where there is uncertainty in the time of our decisions. We distinguish between a priori and a posteriori policies for the recovery of information. An a priori policy is referred to as a policy based on prior knowledge of the environment.

In an a posteriori policy both the prior knowledge and the result of each observation of the environment are used in determining the optimal policy.

In Chapter 4 the a priori information recovery policies are investigated in detail and the conditions for optimum information recovery are found. It is assumed in this chapter that the decision occurs only once. The effect of risk aversion on the optimal information recovery policy is also studied.

Chapter 5 investigate the a posteriori information recovery policies in detail. Necessary conditions for optimum recovery policies are found. Conditions are also found under which the a priori and the a posteriori policies coincide.

Chapter 6 extends the results of Chapter 4 to the case of multiply occurring decisions, namely when the decision may be repeated in time. In this case we often have the opportunity to obtain free information from each decision and use this information for future decisions. The optimality conditions for recovery of information are found under various types of information learned from decisions.

The final chapter summarizes the results of the study and suggests areas for future research.

CHAPTER 2

INFORMATION OUTDATING: PERISHING OR ENHANCEMENT

The purpose of this chapter is to investigate the process of outdating of information in a dynamic, probabilistic environment. We will investigate the manner in which the outdating of information depends on the dynamics of the environment, the decision for which the information is used, and the characteristics of the information itself. The possible patterns of the process under various conditions are studied.

2.1 Formulation of the Problem

The outdating of information refers to the changes over time of the expected value of a piece of information about an uncertain variable at a given time. Formally, let $v_1(\underline{s},\underline{d})$ denote the payoff of a decision D , where $\underline{s}$ is the vector of uncertain states and $\underline{d}$ is the vector of decisions to be set by the decision maker. In a dynamic, probabilistic environment $\underline{s}$ would be changing over time in an uncertain manner. The payoff to the decision maker, if the decision is made at time t , is $v_1(\underline{s}(t),\underline{d})$. Suppose now that the decision maker is offered the perfect information about $\underline{s}(o)$. The value of this information would depend, in general, on time t when the decision maker wants to (or must) make the decision. We refer to this information as unfresh, delayed or old (as opposed to fresh, prompt, or new information). We often expect that the value of this information will be diminishing with t . This is called information perishing [5]. The process of information outdating depends, in general, on the dynamics of $\underline{s}(t)$, on the decision for which the information is used (payoff function v_1), and on the characteristics

of the information itself. The manner in which each of these factors influences the outdating of information is studied in this chapter.

In a more general case, the payoff function v_1 itself may be changing over time. However, since we are primarily interested in the dynamics of information, we exclude that possibility here.

2.2 Notation

The following notation is used throughout the research:

$v_1(\underline{s},\underline{d})$:	Payoff function
$\underline{s} = \underline{s}(t)$:	State vector (stochastic)
$\underline{z}(t) = \eta(\underline{s}(t))$:	Information structure (observation). η may be a noiseless (many-to-one) or noisy (one-to-many) maping. $\underline{z}$ denotes the realization of η over $\underline{s}$.
$\underline{y}(t)$:	Information available at time t . $\underline{y}(t) = \underline{z}(t)$ if information is fresh. $\underline{y}(t) = \underline{z}(t-\tau)$ if information is τ units old (delayed).

2.3 Information Outdating: a priori vs. a posteriori

It is very helpful to distinguish, at the outset, between a priori and a posteriori outdating of information. Both notions are used throughout this research. A priori outdating of information refers to the changes in the a priori expected value of information about $\underline{s}(t-\tau)$ for making the decision at time t . The actual realization of $\underline{s}(t-\tau)$ is not known, when the expected value of the information is

evaluated. This is the case, for instance, when information about the state is available with a delay τ . We may define:

$$V_{\eta(t-\tau)}(t) = \text{expected value of information structure } \eta(\underline{s}(t-\tau)) \text{ for making the decision at time } t \; . \qquad (2.1)$$

We may expect $V_{\eta(t-\tau)}(t)$ to be decreasing as τ increases, (see Fig. 2.1).

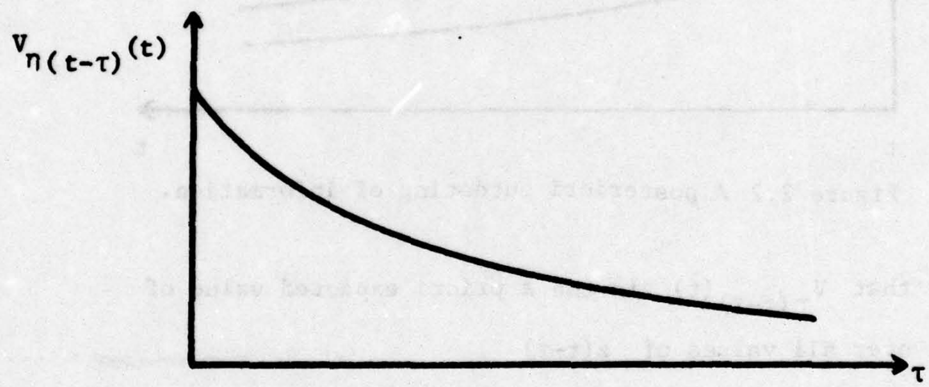

Figure 2.1 A priori outdating of information.

A posteriori outdating of information refers to the case, where the actual realization of the observation is known and we are concerned with the usefulness of this information (data), as time passes. Let us define:

$$V'_{\underline{z}(t_o)}(t) = \text{expected gain at time } t \text{ of using an already known information, namely } \underline{z}(t_o) \; . \qquad (2.2)$$

Note that $V'_{\underline{z}(t_o)}(t)$ depends, in general, on $\underline{z}(t_o)$,(see Fig. 2.2).

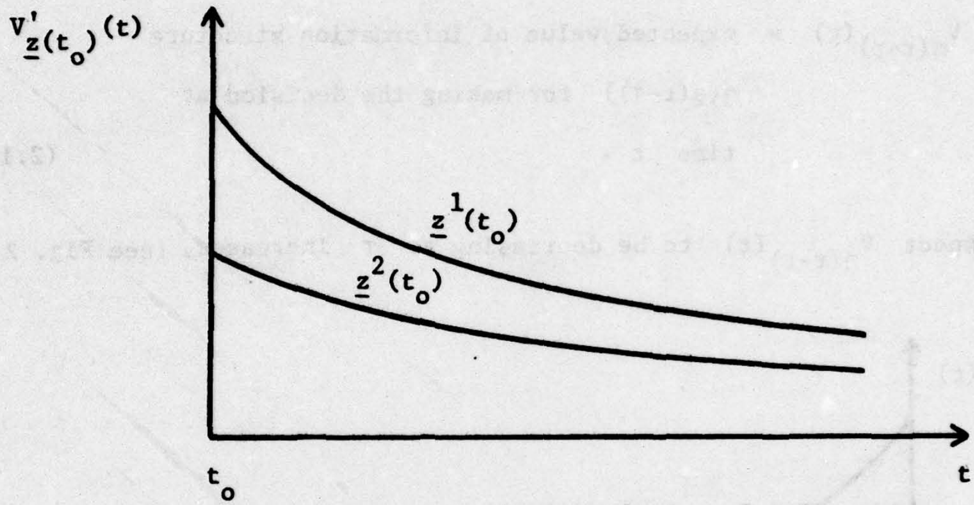

Figure 2.2 A posteriori outdating of information.

It is clear that $V_{\eta(t-\tau)}(t)$ is the a priori expected value of $V'_{\underline{z}(t-\tau)}(t)$ over all values of $\underline{z}(t-\tau)$.

2.4 Information Perishing Rate

Suppose that $V_{\eta(t-\tau)}(t)$ is a decreasing function of τ . This is called "Information Perishing" [5], and the "Rate of Information Perishing" is defined as:

$$\rho = \frac{V_{\eta(t-\tau)}(t)}{V_{\eta(t-\tau-1)}(t)} = \frac{\text{expected value of information with delay } \tau}{\text{expected value of information with delay } \tau+1} \tag{2.3}$$

This definition is useful if ρ is a function of τ but not t . If $\underline{s}(t)$ is stationary and the payoff function v_1 does not change in

time, then ρ is independent of t. Note that using the a posteriori value of information (V') instead of the a priori value of information (V) in (2.3) will not give us a good measure of information perishing because a rate so defined would depend not only on $t - t_o$, but also on $\underline{z}(t_o)$, namely the realization of the observation at time t_o.

2.5 Limitations and Assumptions

One important limitation of dealing with information is that it cannot be measured like other commodities. One cannot assign a meaningful quantity to a piece of information (as a measure of its information content) in such a manner that the information with the larger quantity is always more valuable than the one with the smaller quantity. This is because the value of a piece of information is not independent of the decision for which the information is used. In particular, information structure η_1 may be more valuable than η_2 for making a decision D_1, while η_2 is more valuable than η_1 for decision D_2.* Therefore we have to specify the decision, or equivalently, the payoff function $v_1(\underline{s},\underline{d})$ (and also the utility function when risk aversion is significant). The results would, of course, be limited to the particular payoff function which is used. The more general the payoff function, the more general the results will be. The mathematical complexities, however, often limit our choices of the payoff function. In almost all of this study we assume a quadratic payoff function, which is a good approximation in most practical situations:

---

* Although one can find information structures η_1 and η_2, such that η_1 is always (for every potential user) more valuable than η_2 (η_1 "more informative" than η_2), this relation cannot order all information structures η_1 and η_2.

$$v_1(\underline{s},\underline{d}) = \underline{s}'G\,\underline{d} + \frac{1}{2}\,\underline{d}'H\,\underline{d} \tag{2.4}$$

where $\underline{s}$ and $\underline{d}$ are vectors and G and H are matrices of appropriate dimensions. H is assumed to be negative definite to guarantee the existence of a maximum point. Adding any function of $\underline{s}$ alone to $v_1(\underline{s},\underline{d})$ in (2.4) will change none of our results because it will not change the optimal decision, although the payoff will be different.

The dynamic systems which we study here are mostly linear Markovian. Some of the results, however, are true for any system. The autoregressive systems have also been studied. The systems are assumed to be stationary. We can show, however, that there is no fundamental difficulty in extending the results to nonstationary systems. Another assumption concerns the distribution of $\underline{s}(t)$. It is often assumed that $\underline{s}(t)$ has normal distribution. This assumption simplifies the results and is not an unreasonable assumption because of the additive structue of the dynamic systems studied.

2.6 A Priori Outdating of Information

Throughout the rest of this chapter we will study the a priori outdating of information. In doing so, however, we have to find the a posteriori outdating of information as well. Nevertheless, our main emphasis is on the a priori case. In the following chapters, where we study the recovery of information, the two cases are studied separately.

Assuming a quadratic payoff function

$$v_1(\underline{s},\underline{d}) = \underline{s}'G\,\underline{d} + \frac{1}{2}\,\underline{d}'H\,\underline{d} \tag{2.5}$$

the expected payoff of a decision $\underline{d}$, given information $\underline{y}$ is*

$$\left\langle v_1(\underline{s},\underline{d})|\underline{y},\mathscr{E}\right\rangle = \left\langle \underline{s}'|\underline{y},\mathscr{E}\right\rangle G\,\underline{d} + \frac{1}{2}\underline{d}'H\underline{d} \quad . \tag{2.6}$$

The optimal decision $\hat{d}$ (given information $\underline{y}$) is obtained by setting the derivative of (2.6), with respect to $\underline{d}$, to zero:

$$\frac{\partial}{\partial d}\left\langle v_1(\underline{s},\underline{d})|\underline{y},\mathscr{E}\right\rangle = G\left\langle \underline{s}|\underline{y},\mathscr{E}\right\rangle + H\,\hat{\underline{d}} = 0$$

or

$$\hat{\underline{d}} = -H^{-1}G'\cdot\left\langle \underline{s}|\underline{y},\mathscr{E}\right\rangle = -H^{-1}G'\cdot\tilde{\underline{s}}(\underline{y}) \tag{2.7}$$

$\tilde{\underline{s}}(\underline{y}) = \langle \underline{s}|\underline{y},\mathscr{E}\rangle$ is the posterior mean of $\underline{s}$ given information $\underline{y}$. Substituting $\hat{\underline{d}}$ into (2.6) the maximum expected payoff, given $\underline{y}$, is:

$$\left\langle v_1(\underline{s},\hat{\underline{d}}(\underline{y}))|\underline{y},\mathscr{E}\right\rangle = -\frac{1}{2}\tilde{\underline{s}}'(\underline{y})\cdot G H^{-1}G'\cdot\tilde{\underline{s}}(\underline{y}) \tag{2.8}$$

We assume that $\underline{s}(t)$ is stationary and $\langle \underline{s}(t)|\mathscr{E}\rangle = 0$. The assumption that $\langle \underline{s}(t)|\mathscr{E}\rangle = 0$ is not restrictive because if $\langle \underline{s}(t)|\mathscr{E}\rangle = a \neq 0$, we can always make a change of variable $\underline{s}_1(t) = \underline{s}(t) - a$, where $\underline{s}_1(t)$ has zero mean,and the problem will remain basically the same. By this assumption the maximum expected payoff without information, from (2.8), is

*We use inferential notation which explicitly shows the state of information. $\{x|y,\mathscr{E}\}$ denotes the probability distribution of x given observation y and our prior knowledge $\mathscr{E}$. $\langle x|y,\mathscr{E}\rangle$ and $^v\langle x|y,\mathscr{E}\rangle$ denote the expected value and the variance of x given $(y,\mathscr{E})$, respectively.

zero because $\tilde{\underline{s}}' = \langle \underline{s}(t) | \mathscr{E} \rangle = 0$. Now suppose that our information $\underline{y}(t)$ is the result of observation η with delay τ ,

$$\underline{y}(t) = \underline{z}(t-\tau) = \eta(\underline{s}(t-\tau)) \tag{2.9}$$

From (2.8), the maximum expected payoff with this information is

$$\left\langle v_1(\underline{s}(t), \hat{\underline{d}}(\underline{y}(t)) \,|\, \underline{y}(t) = \eta(\underline{s}(t-\tau)), \mathscr{E} \right\rangle = \tilde{\underline{s}}'(\underline{y}) M \, \tilde{\underline{s}}(\underline{y}) \tag{2.10}$$

where

$$\tilde{\underline{s}}(\underline{y}) = \left\langle \underline{s}(t) \,|\, \underline{y}(t) = \eta(\underline{s}(t-\tau)), \mathscr{E} \right\rangle \tag{2.11}$$

and

$$M = -\frac{1}{2} G H^{-1} G' \tag{2.12}$$

Since the expected payoff with no information is zero, (2.10) gives us the value of information $\underline{y}$, or according to our definition (2.2), the a posteriori value of information $V'_{\underline{z}(t-\tau)}(t)$. Since $s(t)$ is assumed to be stationary, $V'_{\underline{z}(t-\tau)}(t)$ does not depend on t , and we show it by $V'_{\underline{z}(\tau)}$:

$$V'_{\underline{z}(\tau)} = \tilde{\underline{s}}'(\underline{z}(\tau)) \cdot M \cdot \tilde{\underline{s}}(\underline{z}(\tau)) \quad . \tag{2.13}$$

Note that $V'_{\underline{z}(\tau)}$ depends on $\underline{z}$, namely the realization of the observation η at $t-\tau$. A priori value of information structure η with

delay τ ($V_{\eta(t-\tau)}(t)$) can be obtained by taking the expected value of $V'_{\underline{z}(\tau)}$ over all values of $\underline{z}(\tau)$. Again by stationarity of $\underline{s}(t)$, this value does not depend on t and we will show it by $V_{\eta(\tau)}$:

$$V_{\eta(\tau)} = \left\langle V'_{\underline{z}(\tau)} \middle| \mathscr{E} \right\rangle = \left\langle \tilde{\underline{s}}'(\underline{z}(\tau)) \cdot M \cdot \tilde{\underline{s}}(\underline{z}(\tau)) \middle| \mathscr{E} \right\rangle \tag{2.14}$$

For a random vector $\underline{x}$ and a matrix A we have

$$\left\langle \underline{x}' A \underline{x} \middle| \mathscr{E} \right\rangle = \operatorname{trace}\left(A \cdot \sum_{\underline{x}}\right) + \left\langle \underline{x} \middle| \mathscr{E} \right\rangle' A \left\langle \underline{x} \middle| \mathscr{E} \right\rangle \tag{2.15}$$

where $\Sigma_{\underline{x}}$ is the covariance matrix of $\underline{x}$. Therefore (2.14) can be written as:

$$V_{\eta(\tau)} = \operatorname{tr}\left(M \sum_{\tilde{\underline{s}}}\right) + \left\langle \tilde{\underline{s}}(\underline{z}(\tau)) \middle| \mathscr{E} \right\rangle' M \left\langle \tilde{\underline{s}}(\underline{z}(\tau)) \middle| \mathscr{E} \right\rangle$$

but

$$\left\langle \tilde{\underline{s}}(\underline{z}(\tau)) \middle| \mathscr{E} \right\rangle = \left\langle \langle s(t) \,|\, \underline{z}(t-\tau), \mathscr{E} \rangle \middle| \mathscr{E} \right\rangle = \left\langle \underline{s}(t) \middle| \mathscr{E} \right\rangle = 0 \quad .$$

Therefore,

$$V_{\eta(\tau)} = \operatorname{tr}\left(M \cdot \sum_{\tilde{\underline{s}}}\right) \tag{2.16}$$

where

$$\sum_{\tilde{\underline{s}}} = \operatorname{cov}\left\langle \tilde{\underline{s}}(\underline{z}(\tau)) \middle| \mathscr{E} \right\rangle \tag{2.17}$$

Note that $\Sigma_{\tilde{\underline{s}}}$ depends only on τ (and not on $\underline{z}$). Equation (2.16) gives us the expected value of information structure η with delay τ.

The effect of payoff function on the value of information is reflected in matrix M. The dynamics of the system, as well as the characteristics of the information structure and the delay τ, are represented by matrix $\Sigma_{\tilde{\underline{s}}}$ of the covariances of the posterior means. To see the effect of each factor separately, let us first calculate the posterior mean, $\tilde{\underline{s}}(\underline{z}(\tau))$. Expanding over $\underline{s}(t-\tau)$ we have

$$\tilde{\underline{s}}(\underline{z}(\tau)) = \left\langle \underline{s}(t)\,|\,\underline{z}(t-\tau),\mathscr{E} \right\rangle = \int_{\underline{s}(t-\tau)} \left\langle \underline{s}(t)\,|\,\underline{s}(t-\tau),\underline{z}(t-\tau),\mathscr{E} \right\rangle \cdot$$

$$\{\underline{s}(t-\tau)\,|\,\underline{z}(t-\tau),\mathscr{E}\} = \int_{\underline{s}(t-\tau)} \left\langle s(t)\,|\,\underline{s}(t-\tau),\mathscr{E} \right\rangle \{\underline{s}(t-\tau)\,|\,\underline{z}(t-\tau),\mathscr{E}\}$$

(2.18)

$\langle \underline{s}(t)\,|\,\underline{s}(t-\tau),\mathscr{E} \rangle$ is a function of $\underline{s}(t-\tau)$ and τ. Using a linear approximation for this function, and in view of stationarity of $\underline{s}(t)$, we can write:

$$\left\langle \underline{s}(t)\,|\,\underline{s}(t-\tau),\mathscr{E} \right\rangle \cong R(\tau)\cdot\underline{s}(t-\tau) + \underline{q}(\tau)$$

where $R(\tau)$ is a square matrix and $\underline{q}(\tau)$ is a vector. Since $\langle \underline{s}(t)\,|\,\mathscr{E} \rangle = 0$, it follows that $\underline{q}(\tau) = 0$, and we have

$$\left\langle \underline{s}(t)\,|\,\underline{s}(t-\tau),\mathscr{E} \right\rangle \cong R(\tau)\cdot\underline{s}(t-\tau) \quad . \qquad (2.19)$$

We will show later that (2.19) is exact if $\underline{s}(t)$ has normal distribution. Substituting from (2.19) into (2.18) we have

$$\left\langle \underline{s}(t) \,|\, \underline{z}(t-\tau), \mathscr{E} \right\rangle = R(\tau) \cdot \left\langle \underline{s}(t-\tau) \,|\, \underline{z}(t-\tau), \mathscr{E} \right\rangle \tag{2.20}$$

Therefore,

$$\Sigma_{\tilde{\underline{s}}} = \overset{\vee}{\left\langle} \langle \underline{s}(t) \,|\, \underline{z}(t-\tau), \mathscr{E} \rangle \,|\, \mathscr{E} \right\rangle =$$

$$= \overset{\vee}{\left\langle} R(\tau) \cdot \langle \underline{s}(t-\tau) \,|\, \underline{z}(t-\tau), \mathscr{E} \rangle \,|\, \mathscr{E} \right\rangle$$

$$= R(\tau) \cdot \overset{\vee}{\left\langle} \langle \underline{s}(t-\tau) \,|\, \underline{z}(t-\tau), \mathscr{E} \rangle \,|\, \mathscr{E} \right\rangle \cdot R'(\tau)$$

$$= R(\tau) \cdot \Sigma_{\tilde{\underline{s}}_o} \cdot R'(\tau) \tag{2.21}$$

where $\Sigma_{\tilde{\underline{s}}_o}$ is the matrix of covariances of posterior means with fresh information. Note that by stationary assumption, $\Sigma_{\tilde{\underline{s}}_o}$ does not depend on the time of observation $t-\tau$. Equation (2.21) separates the effects of the information structure (η) and the dynamics of the system. Substituting from (2.21) into (2.16) we have

$$V_{\eta(\tau)} \cong \operatorname{tr} \left[M \cdot R(\tau) \cdot \Sigma_{\tilde{\underline{s}}_o} \cdot R'(\tau) \right] \tag{2.22}$$

Equation (2.22) shows how the value of information is related to the decision (payoff function), the dynamics of $\underline{s}(t)$, the information structure η , and the delay τ . The payoff function influences the value of information through matrix $M = -1/2\, G H^{-1} G'$. Information structure η is represented by $\Sigma_{\tilde{\underline{s}}_o}$, namely the matrix of covariance of posterior means with fresh information. The delay τ and the dynamics of the system are reflected in the matrix $R(\tau)$. $R(\tau)$ is the coefficient of the linear approximation of the mean of $\underline{s}(t)$ in terms of $\underline{s}(t-\tau)$. We now show that if $\underline{s}(t)$ has normal distribution, (2.22) is exact and $R(\tau)$ is the "multiple correlation" matrix between $\underline{s}(\tau)$ and $\underline{s}(t-\tau)$. For normal $\underline{s}(t)$, we have:*

$$\left\langle \underline{s}(t)|\underline{s}(t-\tau),\mathscr{E}\right\rangle = \left\langle \underline{s}(t)|\mathscr{E}\right\rangle + \Sigma_{\underline{s}(t),\underline{s}(t-\tau)} \cdot \Sigma^{-1}_{\underline{s}(t-\tau),\underline{s}(t-\tau)}$$

$$\cdot\ [\underline{s}(t-\tau) - \langle \underline{s}(t-\tau)|\mathscr{E}\rangle]$$

Since $\langle s(t)|\mathscr{E}\rangle = \langle s(t-\tau)|\mathscr{E}\rangle = 0$,

$$\left\langle \underline{s}(t)|\underline{s}(t-\tau),\mathscr{E}\right\rangle = \Sigma_{\underline{s}(t),\underline{s}(t-\tau)} \cdot \Sigma^{-1}_{\underline{s}(t-\tau),\underline{s}(t-\tau)} \cdot \underline{s}(t-\tau) .$$

Therefore (2.19) is exact for this case with $R(\tau)$ being

$$R(\tau) = \Sigma_{\underline{s}(t),\underline{s}(t-\tau)} \cdot \Sigma^{-1}_{\underline{s}(t-\tau),\underline{s}(t-\tau)} . \qquad (2.23)$$

---

*See updating relations for multivariate normal distributions in [2] for example.

This matrix is called the multiple correlation matrix between $\underline{s}(t)$ and $\underline{s}(t-\tau)$. If $\underline{s}(t)$ is single, $R(\tau)$ reduces to the correlation coefficient between $s(t)$ and $s(t-\tau)$.

For normal $\underline{s}(t)$ we can find $\Sigma_{\tilde{\underline{s}}_o}$ in terms of the covariance matrices $\Sigma_{\underline{s},\underline{z}}$ and $\Sigma_{\underline{z},\underline{z}}$. The prior variance $\Sigma_{\underline{s}}\left(=\Sigma_{\underline{s},\underline{s}}\right)$ can be decomposed to the expected value of posterior variance $\tilde{\Sigma}_{\underline{s}}$, and the variance of the posterior mean $(\tilde{\underline{s}})$ [16] :

$$\Sigma_{\underline{s}} = \left\langle \tilde{\Sigma}_{\underline{s}} \middle| \mathscr{E} \right\rangle + \Sigma_{\tilde{\underline{s}}}$$

or

$$\Sigma_{\tilde{\underline{s}}} = \Sigma_{\underline{s}} - \left\langle \tilde{\Sigma}_{\underline{s}} \middle| \mathscr{E} \right\rangle \tag{2.24}$$

(2.24) is true in general. For $\underline{s}$ normal we have [2]

$$\tilde{\Sigma}_{\underline{s}} = \Sigma_{\underline{s}} - \Sigma_{\underline{s},\underline{z}} \cdot \Sigma_{\underline{z}}^{-1} \cdot \Sigma_{\underline{z},\underline{s}} \, .$$

Since $\tilde{\Sigma}_{\underline{s}}$ does not depend on $\underline{z}$, from (2.24) we have

$$\Sigma_{\tilde{\underline{s}}} = \Sigma_{\underline{s}} - \tilde{\Sigma}_{\underline{s}} = \Sigma_{\underline{s},\underline{z}} \Sigma_{\underline{z}}^{-1} \Sigma_{\underline{z}.\underline{s}} \, . \tag{2.25}$$

Note that for this case, $\Sigma_{\tilde{\underline{s}}}$ is the difference between the prior and the posterior variance of $\underline{s}$ and can be written in terms of $\Sigma_{\underline{s},\underline{z}}$ and $\Sigma_{\underline{z}}$. Therefore, for a normal state we have:

$$\begin{cases} V_{\eta(\tau)} = \operatorname{tr}\left[M \cdot R(\tau) \cdot \sum_{\tilde{\underline{s}}_o} \cdot R'(\tau)\right] & \text{(exact)} \\ R(\tau) = \sum_{\underline{s}(t),\underline{s}(t-\tau)} \cdot \sum^{-1}_{\underline{s}(t-\tau)} = & \text{multiple correlation between } \underline{s}(t) \text{ and } \underline{s}(t-\tau) \\ \sum_{\tilde{\underline{s}}_o} = \sum_{\underline{s},\underline{z}} \cdot \sum^{-1}_{\underline{z}} \cdot \sum_{\underline{z},\underline{s}} = & \text{reduction in variance as a result of the observation} \end{cases}$$

2.7 Information Outdating in Linear Markovian Systems

The results obtained in Section 2.6 are true for dynamic systems in general. A special case of interest is the linear Markovian system characterized by equation

$$\underline{s}(t) = A\,\underline{s}(t-1) + \underline{\varepsilon}(t) \tag{2.26}$$

where $\underline{s}(t)$ is the vector of states at time t, A is a square matrix and $\underline{\varepsilon}$ is a vector of random "noises." We assume that $\underline{\varepsilon}(t)$, $\underline{\varepsilon}(t-1)$, ... are independent and identically distributed and have zero means (for all t). We also assume that $\underline{s}(t)$ is stationary. This assumption requires that all the eigenvalues of A lie inside the unit circle. Substituting for $\underline{s}(t-1)$ in (2.26) from $\underline{s}(t-1) = A\,\underline{s}(t-2) + \underline{\varepsilon}(t-1)$, we have

$$\underline{s}(t) = A^2\underline{s}(t-1) + A\,\underline{\varepsilon}(t-1) + \underline{\varepsilon}(t)$$

By similar substitutions for $\underline{s}(t-2)$,... $\underline{s}(t-\tau+1)$ we get

$$\underline{s}(t) = A^{\tau}\underline{s}(t-\tau) + A^{\tau-1}\underline{\varepsilon}(t-\tau+1) + A^{\tau-2}\underline{\varepsilon}(t-\tau+2) + \dots + A\,\underline{\varepsilon}(t-1) + \underline{\varepsilon}(t) \tag{2.27}$$

From the independence of the $\underline{\varepsilon}(t)$'s and the stationarity assumption it follows that $\underline{\varepsilon}(t)$ is independent of $\underline{s}(t-1)$, $\underline{s}(t-2)$, ... Therefore from (2.27) we have

$$\left\langle \underline{s}(t) \,|\, \underline{s}(t-\tau), \mathscr{E} \right\rangle = A^{\tau} \cdot \underline{s}(t-\tau) \; . \tag{2.28}$$

Comparing (2.28) with (2.19), $R(\tau)$ for the linear Markovian system of (2.26) is

$$R(\tau) = A^{\tau} \qquad \text{(exact)}$$

Substituting into (2.22) we find the (exact) value of information structure η with delay τ for our Markovian system:

$$V_{\eta(\tau)} = \operatorname{tr}\left[M \cdot A^{\tau} \cdot \sum\nolimits_{\underline{\tilde{s}}_o} \cdot A'^{\,\tau} \right] \tag{2.29}$$

In the following we will investigate the patterns of $V_{\eta(\tau)}$ (as a function of τ) under various circumstances. All the results concern the Markovian system.

<u>Theorem 2.1</u>. If information structure η is perfect, ($\eta(\underline{s}(t)) = \underline{s}(t)$) , then $V_{\eta(\tau)}$ (for Markovian system) always decreases as τ increases (Information Perishing).

<u>Proof</u>: From (2.27) and by the stationarity assumption, $\underline{s}(t)$ can be written as

$$\underline{s}(t) = \underline{\varepsilon}(t) + A\,\underline{\varepsilon}(t-1) + A^2\underline{\varepsilon}(t-2) + \ldots \tag{2.30}$$

From this equation, and the independence of $\varepsilon(t)$'s we have

$$\Sigma_{\underline{s}(t)} = V + AVA' + A^2VA'^2 + \ldots \tag{2.31}$$

where V is the covariance matrix of $\underline{\varepsilon}(t)$. Since information is perfect,

$$\tilde{\underline{s}}_o = \left\langle \underline{s}(t) \mid \eta(\underline{s}(t)), \mathscr{E} \right\rangle = \left\langle \underline{s}(t) \mid \underline{s}(t), \mathscr{E} \right\rangle = \underline{s}(t)$$

Therefore,

$$\Sigma_{\tilde{\underline{s}}_o} = \Sigma_{\underline{s}} = \sum_{i=0}^{\infty} A^i V A'^i \tag{2.32}$$

Substituting into (2.29) we have:

$$V_{\eta(\tau)} = \mathrm{tr}\left[M \cdot A^{\tau} \cdot \left(\sum_{i=0}^{\infty} A^i V A'^i\right) \cdot A'^{\tau}\right]$$

$$= \mathrm{tr}\left[M \cdot \left(\sum_{i=\tau}^{\infty} A^i V A'^i\right)\right]$$

$$= \mathrm{tr}\left(\sum_{i=\tau}^{\infty} MA^i V A'^i\right)$$

$$= \sum_{i=\tau}^{\infty} \mathrm{tr}(MA^i V A'^i)$$

Since M is positive definite (because H is negative definite) and $A^i V A'^i$ is positive definite, $tr(M \cdot A^i V A'^i)$ is positive [3]. It follows, therefore, that $V_{\eta(\tau)}$ decreases as τ increases for all values of τ (Fig. 2.3).

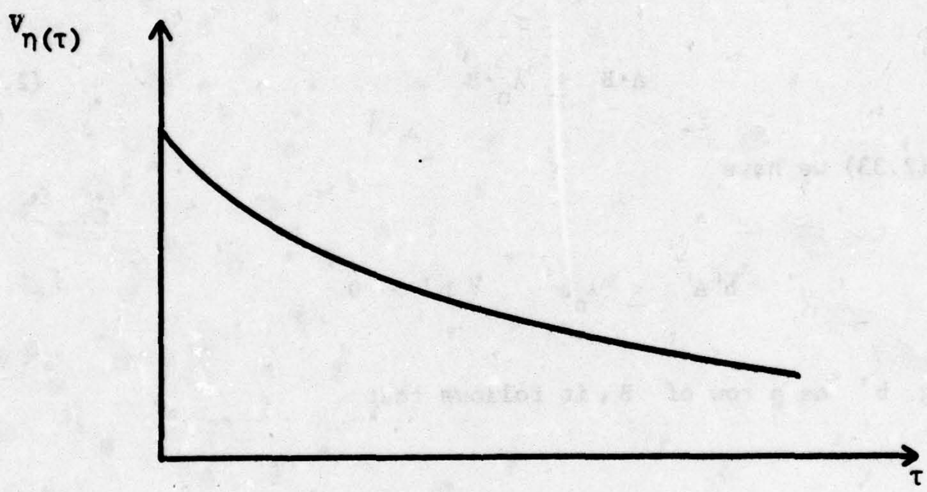

Figure 2.3 Value of perfect information.

Theorem 2.2. If all the matrices A, M, and $\Sigma_{\tilde{s}_o}$ are nonnegative, then $V_{\eta(\tau)}$ is always decreasing with τ and the rate of information perishing $(\rho(\tau))$ is always greater than or equal to $1/\lambda_n^2$, where λ_n is the largest eigenvalue of A.

Proof: We have

$$V_{\eta(\tau+1)} = tr\left(M A^{\tau+1} \sum_{\tilde{s}_o} A'^{\tau+1}\right)$$

$$= tr\left(M A \cdot A^{\tau} \sum_{\tilde{s}_o} A'^{\tau} \cdot A'\right)$$

Let $A^{\tau}\Sigma_{\underline{\tilde{s}}_o} A'^{\tau} = B$, B is nonnegative. Since A is nonnegative if λ_n is the largest eigenvalue of A for any nonnegative vector b we have [3]:

$$A \cdot b \leq \lambda_n \cdot b \tag{2.33}$$

Since all the columns of B are nonnegative it follows that

$$A \cdot B \leq \lambda_n \cdot B \tag{2.34}$$

Also from (2.33) we have

$$b'A' \leq \lambda_n b' \quad \forall\, b' \geq 0$$

Considering b' as a row of B, it follows that

$$B A' \leq \lambda_n B \tag{2.35}$$

From (2.34) and (2.35) we have

$$A B A' \leq \lambda_n^2 \cdot B \tag{2.36}$$

From (2.36), and in view of the nonnegativity of M, we have

$$\begin{aligned} V_{\eta(\tau+1)} &= tr(MABA') \leq tr(M \cdot \lambda_n^2 \cdot B) \\ &= \lambda_n^2 \cdot tr(M \cdot B) \\ &= \lambda_n^2 \cdot V_{\eta(\tau)} \end{aligned}$$

Therefore,

$$\frac{V_{\eta(\tau)}}{V_{\eta(\tau+1)}} \geq \frac{1}{\lambda_n^2} \qquad \forall\, \tau > 0$$

Since $\lambda_n < 1$ (by stationarity assumption), information is always perishing with a rate greater than or equal to $1/\lambda_n^2$.

Theorem 2.3. If $\lambda_1, \lambda_2, \ldots \lambda_n$ are distinct eigenvalues of matrix A, $V_{\eta(\tau)}$ can be written as

$$V_{\eta(\tau)} = \sum_{i,j=1}^{n} c_{ij} (\lambda_i \lambda_j)^{\tau} \tag{2.37}$$

where c_{ij} is constant.

Proof: A can be written as:

$$A = P \Lambda P^{-1}$$

where

$$\Lambda = \begin{bmatrix} \lambda_1 & & 0 \\ & \lambda_2 & \\ & & \ddots \\ 0 & & \lambda_n \end{bmatrix}$$

From (2.29) we have:

$$V_{\eta(\tau)} = \mathrm{tr}\left(M \cdot P \, \Lambda^{\tau} P^{-1} \Sigma_{\tilde{s}_o} \cdot P'^{-1} \Lambda^{\tau} P' \right)$$

Since $tr(X \cdot Y) = tr(Y \cdot X)$ for arbitrary matrices X and Y, we can write

$$V_{\eta(\tau)} = tr\left(P' M P \cdot \Lambda^{\tau} \cdot P^{-1} \sum_{\tilde{s}_o} P'^{-1} \cdot \Lambda^{\tau}\right)$$

Letting $T = P' M P$ and $Q = P^{-1} \Sigma_{\tilde{s}_o} P'^{-1}$, we have

$$V_{\eta(\tau)} = tr\,(T \Lambda^{\tau} Q \Lambda^{\tau})$$

Writing $T = [t_{ij}]$ and $Q = [q_{ij}]$ we have

$$\Lambda^{\tau} Q \Lambda^{\tau} = \left[q_{ij} \lambda_i^{\tau} \lambda_j^{\tau} \right]$$

and

$$V_{\eta(\tau)} = tr(T \cdot \Lambda^{\tau} Q \Lambda^{\tau}) = \sum_{i=1}^{n} \sum_{j=1}^{n} t_{ij} \cdot q_{ji} \lambda_j^{\tau} \lambda_i^{\tau}$$

$$= \sum_{i,j=1} t_{ij} \cdot q_{ji} \cdot (\lambda_i \lambda_j)^{\tau}$$

which is the desired result. Note that

$$c_{ij} = t_{ij} \cdot q_{ji}$$

or

$$C = [c_{ij}] = T \square Q' = T \square Q$$

Therefore the matrix C of coefficients is the congruent product of matrices T and Q .

From Eq. (2.37) we see that $V_{\eta(\tau)}$ is the weighted sum of the geometric terms (in τ) with the base of each term being the product of two eigenvalues of matrix A . This indicates that the eigenvalues of A play an important role in determining the behavior of $V_{\eta(\tau)}$ (as a function of τ) . The effects of the payoff function and the information structure are reflected in the constant coefficients, c_{ij} . In the following we will explore the important modes of behavior of $V_{\eta(\tau)}$.

<u>Remark 2.1</u>. If A is diagonal, and either M or $\Sigma_{\tilde{s}_o}$ is also diagonal, then

$$V_{\eta(\tau)} = \sum_{i=1}^{n} c_{ii} \cdot \lambda_i^{2\tau} \tag{2.38}$$

This corresponds to the case, where either no interaction among information components exists (A and $\Sigma_{\tilde{s}_o}$ diagonal) , or if such an interaction exists it is eliminated by the special form of the payoff function (M diagonal). For this case it is easy to show that

$$\frac{1}{\lambda_n^2} \leq \rho(\tau) = \frac{V_{\eta(\tau)}}{V_{\eta(\tau+1)}} \leq \frac{1}{\lambda_1^2} \, , \qquad \forall\ \eta, \tau \tag{2.39}$$

and

$$\rho(\tau) \geq \rho(\tau+1) \tag{2.40}$$

where λ_1 and λ_n are the smallest and the largest eigenvalues of A, respectively. The intuitive interpretation of (2.39) and (2.40) will be clear if we note that the rate of information perishing for the case of a single variable system,

$$s(t) = \lambda s(t-1) + \varepsilon(t)$$

is $1/\lambda^2$. Since, in this case, no interaction among the variables (or information components) exists, the rate of information perishing is between the rate of information perishing for the fastest perishing component, $(1/\lambda_1^2)$ and the rate for the slowest perishing component, $(1/\lambda_n^2)$. As time passes, the components corresponding to larger perishing rates perish more and the share of slower perishing components in $\rho(\tau)$ becomes larger. This causes $\rho(\tau)$ to decrease as τ increases. For very large τ, the component corresponding to the slowest perishing rate $(1/\lambda_n^2)$ becomes dominant and $\rho(\tau)$ approaches $1/\lambda_n^2$ (See Fig. 2.4).

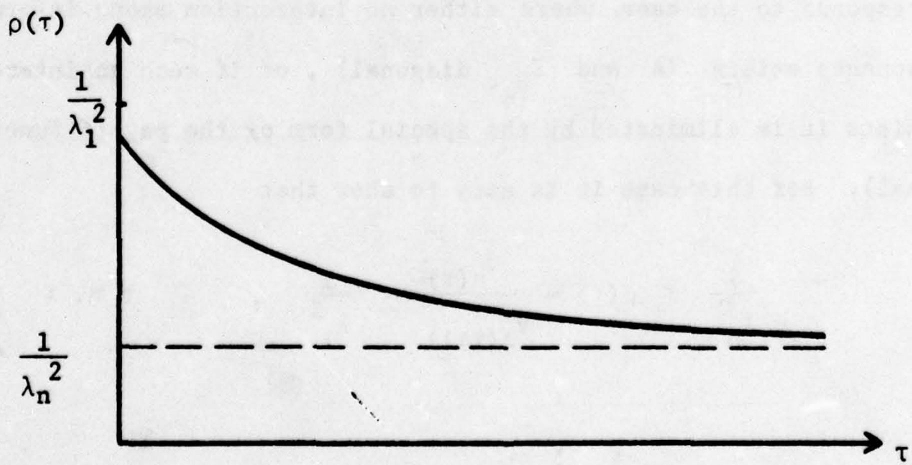

Figure 2.4. $\rho(\tau)$ for diagonal A and M (or A and $\Sigma_{\underset{\sim}{s}_o}$).

Remark 2.2. If η is not perfect, $V_{\eta(\tau)}$ may be increasing with τ in some intervals. Two examples of such cases are given below.

Example 2.1. Let $A = \begin{bmatrix} .2 & 0 \\ 0 & .9 \end{bmatrix}$ and $M = \begin{bmatrix} 1/2 & -1 \\ -1 & 5/2 \end{bmatrix}$.

M is positive definite and eigenvalues of A are less than 1. The choice of $\Sigma_{\tilde{\underline{s}}_o}$ may not be totally arbitrary since $\Sigma_{\tilde{\underline{s}}_o}$ may depend on A (because $\Sigma_{\underline{s}}$ depends on A). However, by appropriate choice of $\Sigma_{\underline{\varepsilon}}$ we have a large range from which to choose $\Sigma_{\tilde{\underline{s}}_o}$. From (2.31),

$$\Sigma_{\underline{s}} = \Sigma_{\underline{\varepsilon}} + A \Sigma_{\underline{\varepsilon}} A' + A^2 \Sigma_{\underline{\varepsilon}} A'^2 + \cdots$$

Letting

$$\Sigma_{\underline{\varepsilon}} = \begin{bmatrix} v_1 & 0 \\ 0 & v_2 \end{bmatrix}$$

we find

$$\Sigma_{\underline{s}} = \begin{bmatrix} \dfrac{v_1}{1-a_{11}^2} & 0 \\ 0 & \dfrac{v_2}{1-a_{22}^2} \end{bmatrix}$$

where $a_{11} = .2$ and $a_{22} = .9$ are elements of matrix A. Since the choices of v_1 and v_2 are arbitrary (subject to being positive), σ_{s_1} and σ_{s_2} can, in fact, be chosen arbitrarily. Let $\Sigma_{\underline{s}}$ be

$$\Sigma_{\underline{s}} = \begin{bmatrix} 4 & 0 \\ 0 & 2 \end{bmatrix} .$$

To find $\Sigma_{\underline{\tilde{s}}}$ we need to have the covariance matrix of $(\underline{s},\underline{z})$. Let this matrix be:

$$\begin{array}{c} \\ s_1 \\ s_2 \\ z_1 \\ z_2 \end{array} \begin{array}{c} \begin{array}{cccc} s_1 & s_2 & z_1 & z_2 \end{array} \\ \begin{bmatrix} 4 & 0 & \sqrt{2} & 0 \\ 0 & 2 & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ \sqrt{2} & \frac{1}{\sqrt{2}} & 1 & 0 \\ 0 & \frac{1}{\sqrt{2}} & 0 & 1 \end{bmatrix} \end{array}$$

The only condition on this matrix is that it must be positive semidefinite. This condition is satisfied. Assuming that s and z have normal distribution, $\Sigma_{\underline{\tilde{s}}}$ can be calculated from (2.25):

$$\Sigma_{\underline{\tilde{s}}_o} = \Sigma_{\underline{s},\underline{z}} \cdot \Sigma_{\underline{z}}^{-1} \cdot \Sigma_{\underline{z},\underline{s}}$$

$$= \begin{bmatrix} \sqrt{2} & 0 \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \sqrt{2} & \frac{1}{\sqrt{2}} \\ 0 & \frac{1}{\sqrt{2}} \end{bmatrix} = \begin{bmatrix} 2 & 1 \\ 1 & 1 \end{bmatrix}$$

Using this matrix we can now calculate the value of information. From theorem (2.3) we have

$$V_{\eta(\tau)} = \sum_{i,j=1}^{2} c_{ij}(\lambda_i\lambda_j)^{\tau}$$

where c_{ij} is the $(ij)^{th}$ element of the matrix C given by:

$$C = T \square Q = M \square \Sigma_{\tilde{s}_o} = \begin{bmatrix} \frac{1}{2} & -1 \\ -1 & \frac{5}{2} \end{bmatrix} \begin{bmatrix} 2 & 1 \\ 1 & 1 \end{bmatrix} = \begin{bmatrix} 1 & -1 \\ -1 & \frac{5}{2} \end{bmatrix}$$

Eigenvalues of A are $\lambda_1 = .2$ and $\lambda_2 = .9$; therefore,

$$V_{\eta(\tau)} = (.04)^{\tau} - 2(.18)^{\tau} + 2.5(.81)^{\tau}$$

$V_{\eta(\tau)}$ as a function of τ is shown in Fig. 2.5.

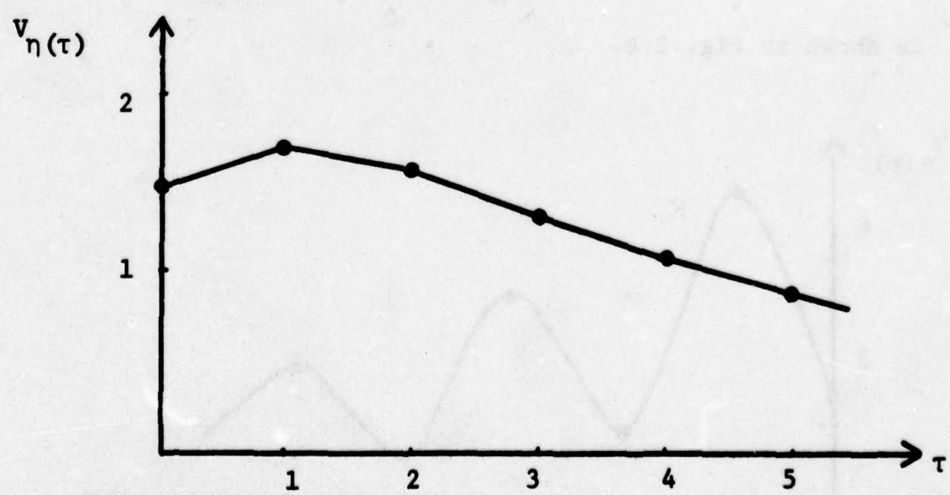

Figure 2.5 $V_{\eta(\tau)}$ as a function of τ in Example 2.1 .

Notice that the value of information with one unit delay (and even two units delay) is greater than the value of fresh information. A more interesting case is given in Example 2.2 below.

Example 2.2 Let M and $\Sigma_{\tilde{\underline{s}}_o}$ be the same as in Example 1, but $A = \begin{bmatrix} .9 & 0 \\ 0 & -.9 \end{bmatrix}$. Notice that the same $\Sigma_{\tilde{\underline{s}}_o}$ as in Example 1 can be chosen with the new matrix A, because $\Sigma_{\underline{s}}$ could be chosen independently of a_{11} and a_{22}. For this case we have:

$$C = T \,\square\, Q = M \,\square \sum_{\tilde{\underline{s}}_o} = \begin{bmatrix} 1 & -1 \\ -1 & \frac{5}{2} \end{bmatrix}$$

and

$$V_{\eta(\tau)} = 3.5(.81)^{\tau} - 2(-.81)^{\tau}$$

$V_{\eta(\tau)}$ is shown in Fig. 2.6.

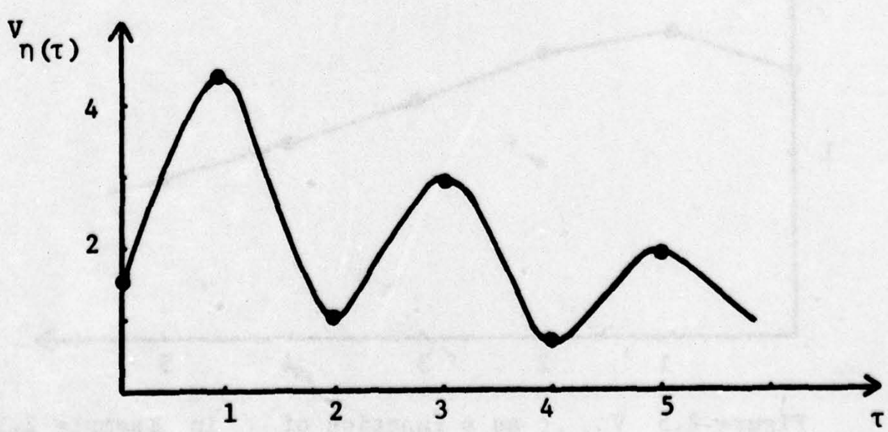

Figure 2.6 $V_{\eta(\tau)}$ in Example 2.2.

Remark 2.3. It is easy to show that if some of the eigenvalues of A are complex, $V_{\eta(\tau)}$ will include, in general, terms of the form $\beta^{\tau}\sin^{2}(\theta\tau+\alpha)$, that is, damped oscillatory terms.

We have seen in the above examples that the value of information does not always perish, as it becomes outdated. This result is rather counter-intuitive considering that our systems have been Markovian. We will give an intuitive interpretation for this mode of behavior of value of information in Section 2.9, where we study the autoregressive systems. The examples also show that the value of information may have various patterns which depend, to a large extent, on the structure of the dynamic system (eigenvalues of matrix A). It must be noted, however, that other factors, namely the payoff function and the information structure itself, can also have important influences on the behavior of the value of information. We have shown, for example, that if the information is perfect, then it will always be perishing, regardless of matrices A or M.

The enhancement of information with delay has an interesting implication for timing of the information acquisition. Suppose that the time of the decision is fixed. We can obtain information about the state of the system at a point in the past such that the expected value of this information is maximum at the time of the decision. If the timing of the decision is flexible, however, we can buy information and choose the time of the decision according to the information, such that the payoff is maximized.

2.8 Information Outdating in the Bidding Example

In this section we will study the outdating of information in the bidding example of Chapter 1. For the bidding model itself we will use the model introduced by Howard [6]. According to this model, the lowest competitive bid (ℓ) and the cost (to us) of performing the contract (p) are uncertain. The optimum bid must be decided based on information about the uncertain variables. There is no uncertainty in the time of the bidding in Howard's model. We assume, however, that the time of the bidding is uncertain and that our uncertain variables change over time. As a result, our information about these variables becomes outdated by passage of time.

Let b denote our bid, the profit v is

$$v(b,\ell,p) = \begin{cases} b - p & b < \ell \\ 0 & b \geq \ell \end{cases}$$

It is easy to show that the expected profit of bidding b is

$$\langle v(b,\ell,p)|b,y,\mathscr{E}\rangle = \{\ell > b|y,\mathscr{E}\} \cdot [b - \langle p|y,\mathscr{E}\rangle] \qquad (2.41)$$

where y denotes our information about the uncertain variables at the bidding time. We assume, for simplicity, that our knowledge about the lowest bid (ℓ) does not change over time and we assign a uniform distribution to this variable (see Fig. 2.7).

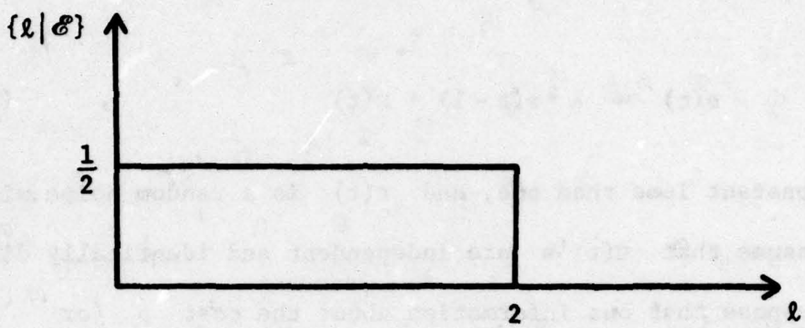

Figure 2.7. Probability distribution of the lowest competitive bid.

The expected profit from (2.41) reduces to

$$\langle v(b,\ell,p)\,|\,b,y,\mathscr{E}\rangle \;=\; \frac{1}{2}(2-b)\cdot[b-\langle p|y,\mathscr{E}\rangle] \tag{2.42}$$

The bid which maximizes the expected profit can be obtained by setting the derivative of (2.42) with respect to b to zero. We find

$$\hat{b} \;=\; 1+\frac{1}{2}\langle p|y,\mathscr{E}\rangle \tag{2.43}$$

and the maximum expected profit is

$$V_1(y) \;=\; \langle v(\hat{b},\ell,p)\,|\,y,\mathscr{E}\rangle$$

$$=\; \frac{1}{2}\left(1-\frac{1}{2}\langle p|y,\mathscr{E}\rangle\right)^2 \tag{2.44}$$

We assume, as was suggested in Chapter 1, that from our past experience we know that the cost of performing the contract (p) has a constant mean

over time (m), and that its variation from its mean ($s = \Delta p$) changes over time, according to the linear Markovian model of

$$s(t) = \lambda \cdot s(t-1) + \varepsilon(t) \tag{2.45}$$

where λ is a constant less than one, and $\varepsilon(t)$ is a random noise with zero mean. We assume that $\varepsilon(t)$'s are independent and identically distributed. Now suppose that our information about the cost p (or equivalently s) at time t is the cost at time zero

$$y(t) = s(o) = s_o \quad .$$

From (2.44) the expected profit of bidding at time t with this information is

$$\begin{aligned} V_1(t,s_o) &= \langle v(\hat{b}(t), \ell, p(t) | s_o, \mathscr{E} \rangle \\ &= \frac{1}{2}[1 - \frac{1}{2}\langle p(t)|s_o,\mathscr{E}\rangle]^2 \\ &= \frac{1}{2}[1 - \frac{1}{2}(m + \langle s(t)|s_o,\mathscr{E}\rangle]^2 \end{aligned} \tag{2.46}$$

but from (2.45) it is easy to show that

$$\langle s(t)|s_o,\mathscr{E}\rangle = \lambda^t s_o \tag{2.47}$$

Therefore, we have

$$V_1(t,s_o) = \frac{1}{2}\left(1 - \frac{m}{2} - \frac{1}{2}\lambda^t s_o\right)^2 \tag{2.48}$$

$V_1(t,s_o)$, namely the maximum expected profit of the bidding at time t with information s_o , is shown in Fig. 2.8. For $s_o < 0$, the expected

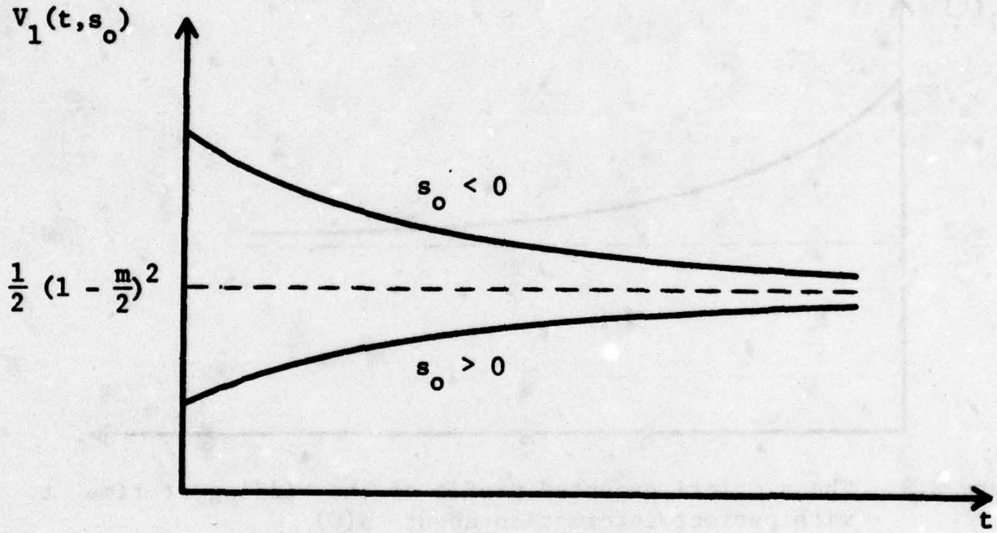

Figure 2.8. Expected profit of the bidding time at time t , given s_o .

profit is decreasing with t . It is increasing for $s_o > 0$. Intuitively, when $s_o < 0$, the cost of performing the contract at time zero is lower than the average cost m , but from Eq. (2.45) this advantage tends to fade out in time. Therefore, the expected profit is decreasing in time. The converse is true, when $s_o > 0$. The a priori expected profit at time t is

$$\bar{V}_1(t) = \langle V_1(t,s_o) | \mathscr{E} \rangle = \frac{1}{2}\left[(1 - \frac{m}{2})^2 + \frac{1}{4}\sigma_s \lambda^{2t} \right] \tag{2.49}$$

where σ_s is the variance of s . $\bar{V}_1(t)$ is shown in Fig. 2.9. It is easy to show that the expected profit of the bidding with no information

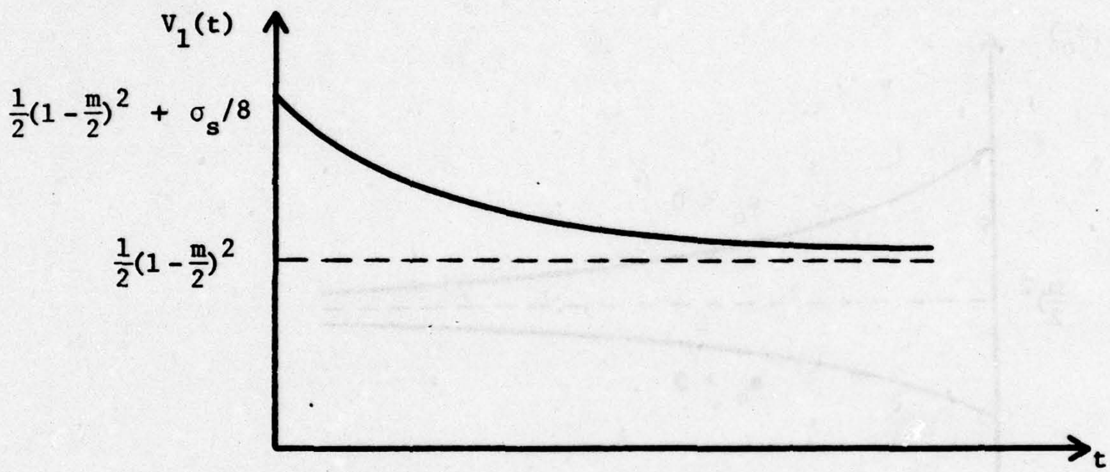

Figure 2.9. The a priori expected profit of the bidding at time t with perfect information about $s(0)$.

is $\frac{1}{2}(1 - \frac{m}{2})^2$. Therefore, the a priori value of perfect information about $s(0)$ at time t is $\frac{1}{8}\sigma_s\lambda^{2t}$. Consequently the information is always perishing and the rate of information perishing is $1/\lambda^2$. If our observation of $s(0)$ is not perfect, then σ_s in (2.49) will be replaced by $\sigma_{\tilde{s}}$, namely the variance of the posteriori mean of $s(0)$, given the observation. We notice that the rate of information perishing remains the same as the case of the perfect information.

2.9 Information Outdating in Autoregressive Systems

A process which is used frequently for modeling the real-world processes is the autoregressive process [4]. The autoregressive process of order p is defined by the equation

$$s(t) = \alpha_1 s(t-1) + \alpha_2 s(t-2) + \cdots + \alpha_p s(t-p) + \varepsilon(t) \quad (2.50)$$

The state of the system at each time depends on its states at the last p points of time, and a random noise $\varepsilon(t)$. We assume that $\varepsilon(t)$'s are independent and identically distributed and $s(t)$ is stationary.*
For simplicity, we assume that $s(t)$ is a single variable. Equation (2.50) can also be written as a linear Markovian system in matrix form:

$$\underbrace{\begin{bmatrix} s(t) \\ s(t-1) \\ \vdots \\ s(t-p+1) \end{bmatrix}}_{\underline{s}(t)} = \underbrace{\begin{bmatrix} \alpha_1 & \alpha_2 & \cdots & \alpha_p \\ 1 & 0 & \cdots & 0 \\ 0 & 1 & 0 & 0 \\ \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & \cdots 1 & 0 \end{bmatrix}}_{A} \cdot \underbrace{\begin{bmatrix} s(t-1) \\ s(t-2) \\ \vdots \\ s(t-p) \end{bmatrix}}_{\underline{s}(t-1)} + \underbrace{\begin{bmatrix} \varepsilon(t) \\ 0 \\ \vdots \\ 0 \end{bmatrix}}_{\underline{\varepsilon}(t)} \quad (2.51)$$

The first row in (2.51) is the same as equation (2.50) and the rest of the rows are identities $s(t-i) = s(t-i)$ for $i = 1,2, \ldots p-1$.

---

*Condition for the stationarity of $s(t)$ is that all the roots of the characteristic polynomial

$$x^p - \alpha_1 x^{p-1} - \alpha_2 x^{p-2} - \cdots - \alpha_p = 0$$

lie inside the unit circle.

Since this system is Markovian, according to Theorem 2.1 the value of perfect information about $\underline{s}(t-\tau)$ (that is perfect information about p consecutive states of the original system: $s(t-\tau)$, $s(t-\tau-1)$, ... , $s(t-\tau-p+1)$) always decreases, as delay τ increases. We will see, however, that the value of perfect information may be increasing with delay, if the information does not include all the p consecutive states in the past. We are interested, in particular, in the value of information about the state at only one point of time in the past. Assuming a quadratic payoff function in a single variable, $s(t)$,

$$v(s,d) = g \cdot s(t) \cdot d + \frac{1}{2} h \cdot d^2 \tag{2.52}$$

the value of information structure η with delay τ from Eq. (2.22) reduces to:

$$V_{\eta(\tau)} \cong m \cdot \sigma_{\tilde{s}_o} \cdot r_\tau^2 \tag{2.53}$$

where $m = 1/2 \cdot g^2/h$ is a constant, $\sigma_{\tilde{s}_o} = {}^{V}\langle\langle s(t)|\eta(s(t)),\mathscr{E}\rangle|\mathscr{E}\rangle$ is the variance of the posterior mean with fresh information and is constant by the stationarity assumption, and r_τ is the coefficient of the linear approximation of the mean of $s(t)$ as a function of $s(t-\tau)$. If $s(t)$ has normal distribution (which is typically the case because of the additive form of the process 2.50), (2.53) is exact and r_τ is the correlation coefficient between $s(t)$ and $s(t-\tau)$. Note that the behavior of $V_{\eta(\tau)}$ depends only on the behavior of r_τ . In

particular, and in constrast to the vector case, there is no difference in the behavior of $V_{\eta(\tau)}$ whether information is perfect or not.

If we multiply both sides of (2.50) by $s(t-\tau)$, take the expected value, and divide both sides by the variance of $s(t)$, we get

$$r_\tau = \alpha_1 r_{\tau-1} + \alpha_2 r_{\tau-2} + \cdots + \alpha_p r_{\tau-p} \tag{2.54}$$

Note that r_τ satisfies the same difference equation as the equation for $s(t)$ with $\varepsilon(t) = 0$. r_τ is therefore the same as the homogeneous solution of $s(t)$, and the value of information is a constant times the square of this homogeneous solution. The solution to the difference equation (2.54), can be written as

$$r_\tau = \sum_{i=1}^{p} c_i \lambda_i^\tau \tag{2.55}$$

where c_i is constant and λ_i is the i^{th} root of the characteristic equation of (2.54) (or equivalently the eigenvalues of matrix A in (2.51)). In the following example we will study the autoregressive process of the second order (AR(2)) in more detail

<u>Example 2.3</u>. Information outdating in AR(2): for AR(2) we have

$$s(t) = \alpha_1 s(t-1) + \alpha_2 s(t-2) + \varepsilon(t) \tag{2.56}$$

The roots of the characteristic equation of this system are:

$$\lambda_1 = \frac{\alpha_1}{2} + \sqrt{\left(\frac{\alpha_1}{2}\right)^2 + \alpha_2}$$

$$\lambda_2 = \frac{\alpha_1}{2} - \sqrt{\left(\frac{\alpha_1}{2}\right)^2 + \alpha_2}$$

To have a stationary process, λ_1 and λ_2 must lie inside the unit circle. This condition gives us the following conditions on α_1 and α_2 :

$$\begin{cases} \alpha_1 + \alpha_2 < 1 \\ \alpha_2 - \alpha_1 < 1 \\ -1 < \alpha_2 < 1 \end{cases} \tag{2.57}$$

Points (α_1, α_2), satisfying conditions (2.57), lie inside the triangle of Fig. 2.10. In the shaded region of the triangle $\alpha_1^2 + 4\alpha_2 < 0$, and

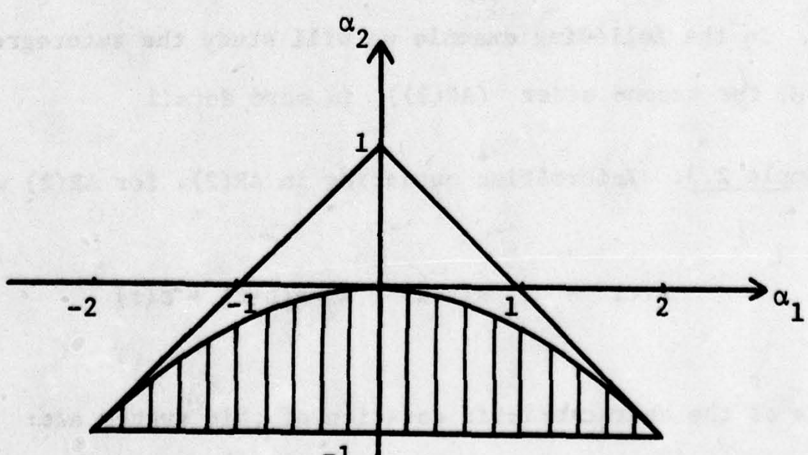

Figure 2.10 Stationary triangle for AR(2) .

the roots of the characteristic equation are complex. For a stationary process the following observations are made (proofs are simple and have been omitted):

$$(1) \qquad V_{\eta(\tau)}(\alpha_1,\alpha_2) = V_{\eta(\tau)}(-\alpha_1,\alpha_2) , \qquad \forall\ \eta\ ,\ \tau$$

that is, the value of any information structure η with any delay τ is the same for two processes with the same α_2 but the two α_1 being negative of each other. In other words, $V_{\eta(\tau)}$ is symmetric with respect to the axis α_2 in Fig. 2.10. Note that in the process with $\alpha_1 > 0$, s(t) changes much slower than does s(t) in the process with $-\alpha_1 < 0$ (Fig.2.11), but the value of information changes at the same rate for both processes.

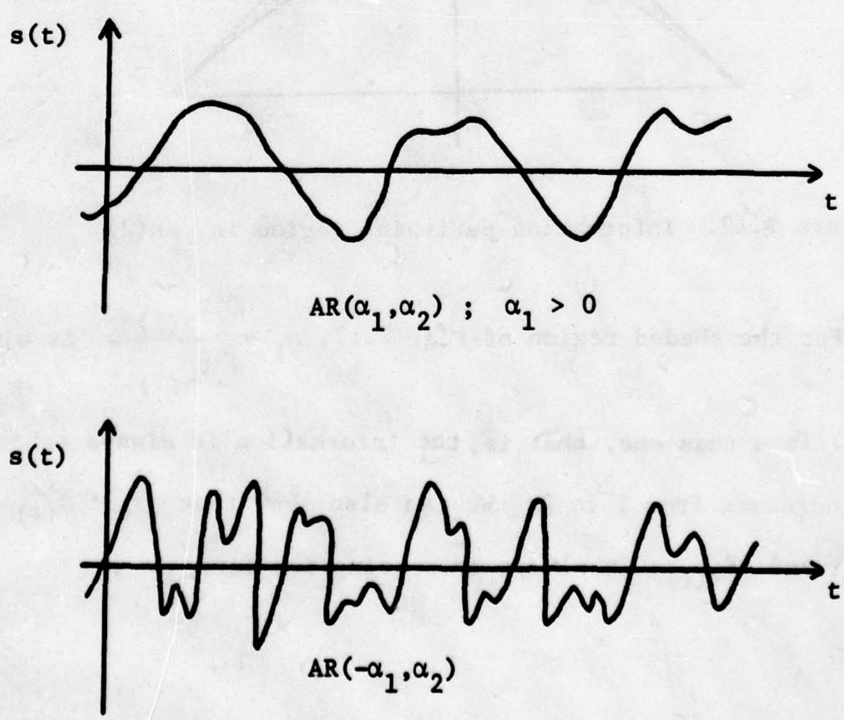

Figure 2.11 AR(2) processes with opposite sign for α_1 .

(2) $V_{\eta(\tau)}(\alpha_1,\alpha_2) \geq V_{\eta(\tau)}(\alpha_1,-\alpha_2)$ for $\alpha_2 > 0$, for all η and τ ; that is, for two processes with the same α_1 , and with the two α_2 being negative of each other, the value of any information structure with any delay will be greater for the process with positive α_2 .

(3) $V_{\eta(\tau)}(\alpha_1,\alpha_2) \geq V_{\eta(\tau+1)}(\alpha_1,\alpha_2)$ for all η and τ , only if (α_1,α_2) lie in the shaded region of Fig. 2.12.

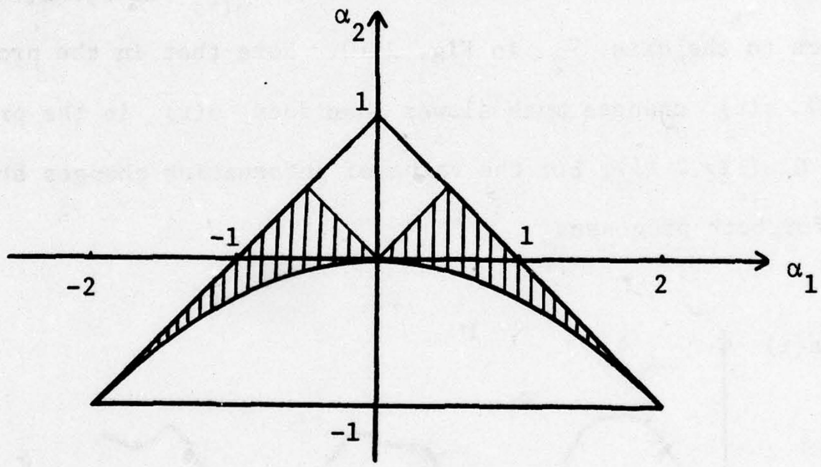

Figure 2.12. Information perishing region in AR(2)

(4) For the shaded region of Fig. 2.13, $\rho_1 = \dfrac{V_{\eta(1)}}{V_{\eta(2)}}$, is always (for all η) less than one, that is, the information is always enhanced if the delay increases from 1 to 2. We can also show that $\rho_\tau \cdot \rho_{\tau+1} \geq 1$ for all τ , and $V_{\eta(\tau)}$ is always decreasing for large τ .

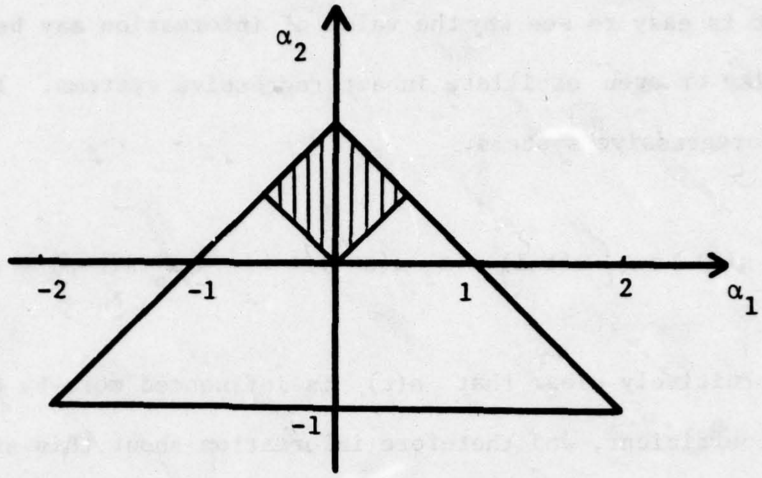

Figure 2.13. Information enhancing region (from $\tau = 1$ to $\tau = 2$) in AR(2) .

The result that the value of information with two units of delay is always more than the value of information with one unit of delay is easy to interpret in this case. Note that in the shaded region of Fig. 2.13, $\alpha_2 > |\alpha_1|$. Therefore, from (2.56) s(t) is determined more by s(t-2) than by s(t-1), and thus it is more valuable to know s(t-2) than s(t-1) .

(5) In the shaded region of Fig. 2.10 where λ_1 and λ_2 are complex, r_τ is a damped oscillatory function of τ . $V_{\eta(\tau)}$, therefore, has the shape of Fig. 2.14.

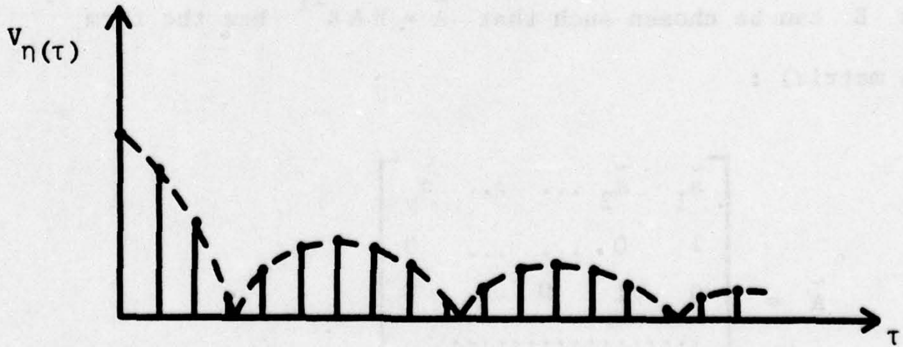

Figure 2.14. $V_{\eta(\tau)}$ in AR(2) with complex characteristic roots.

It is easy to see why the value of information may be increasing with delay or even oscillate in autoregressive systems. In the equation for autoregressive systems,

$$s(t) = \alpha_1 s(t-1) + \alpha_2 s(t-2) + \cdots + \alpha_p s(t-p) + \varepsilon(t)$$

it is intuitively clear that $s(t)$ is influenced more by a state with a larger coefficient, and therefore information about this state can be more valuable than the information about a more recent state. We did not have this clear interpretation for the Markovian system. However, since we can transform a Markovian system into a form equivalent to an autoregressive system, the results for the Markovian system can be interpreted in the same manner. Consider the Markovian system

$$\underline{s}(t) = A\, \underline{s}(t-1) + \underline{\varepsilon}(t)$$

By a transformation $\underline{x}(t) = B\, \underline{s}(t)$ where B is an invertable matrix, we have

$$\underline{x}(t) = B A B^{-1}\, \underline{x}(t-1) + B\, \underline{\varepsilon}(t) \qquad (2.58)$$

The matrix B can be chosen such that $\tilde{A} = B A B^{-1}$ has the form (companion matrix) :

$$\tilde{A} = \begin{bmatrix} \tilde{a}_1 & \tilde{a}_2 \ldots & \ldots & \tilde{a}_p \\ 1 & 0 \ldots & \ldots & 0 \\ 0 & 1 & 0 \ldots & 0 \\ \ldots\ldots & \ldots\ldots & \ldots\ldots & \\ 0 & \ldots\ldots & 1 & 0 \end{bmatrix}$$

Comparing with Eq. (2.51) we notice that the new system represents an autoregressive system of order p (the dimension of A) :

$$x_1(t) = \tilde{a}_1 x_1(t-1) + \tilde{a}_2 x_1(t-2) + \cdots + \tilde{a}_p x_1(t-p) + \tilde{\varepsilon}(t) \qquad (2.59)$$

Our information in the new system is about the vecotr $[x_1(t), x_1(t-1), \ldots x_1(t-p+1)]$ and in view of (2.59) it is easy to see why the value of information may be increasing with delay.

2.10 Comparative Values of Delayed Information Structures

In this section we will explore the changes in the relative values of information structures as they become old. Consider the following question: Information structure η is more valuable than η' when both η and η' are fresh. Is η more valuable than η' when both have a delay τ ? It is clear that the answer to this question may depend not only on η and η' (and τ), but also on the decision for which the information is used, as well as on the properties of the dynamic system. We can think of cases for which η is always perishing but η' may be enhancing in some interval, such that $\eta'(\tau)$ is more valuable than $\eta(\tau)$ (See Fig. 2.15). If however, η and η' are such that the

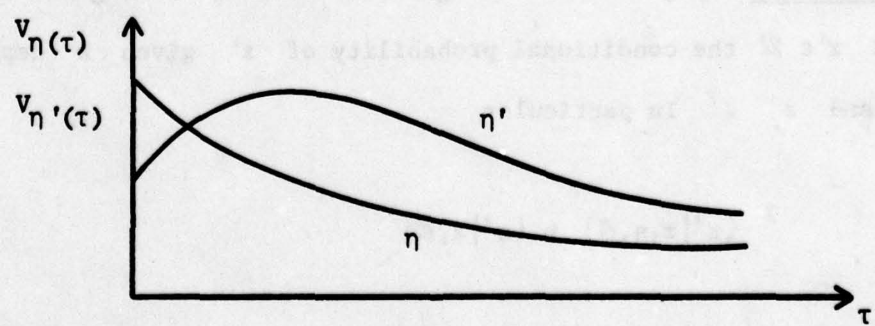

Figure 2.15 Comparative values of delayed information structures.

superiority of η over η' is not limited to a particular decision, then such a property may be preserved or reduced to weaker but similar properties, as η and η' are delayed. Some comparison of information structures can be made according to the following definitions [9] :

Definition (1). $Z \; f \; Z'$: set Z of observation z $(z=\eta(s))$ is finer than set Z' of observations $z'(z'=\eta'(s))$, Fig. 2.16 (η and η' must be many-to-one or so called "noiseless" mappings).

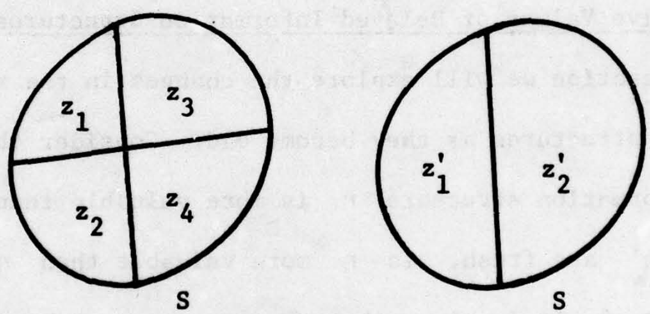

Figure 2.16 Set Z finer than set Z' .

As shown in Fig. 2.16, the partition of the set S (of state variable s) by η is finer than the partition by η' .

Definition (2). $Z \; g \; Z'$: Z is "garbled" into Z' . $Z \; g \; Z'$ if for all $z \in Z$ and $z' \in Z'$ the conditional probability of z' given z depends only on z and z' . In particular

$$\{z'|z,s,\mathscr{E}\} = \{z'|z,\mathscr{E}\}$$

that is, if z is known, the distribution of z' does not depend on the state s itself. This is also called "cascaded information," (See Fig. 2.17)

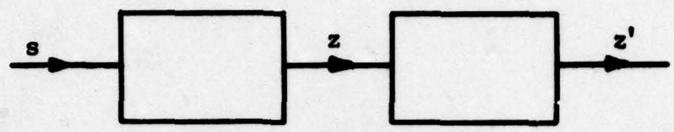

Figure 2.17 Cascaded information.

<u>Definition (3)</u>. $\eta \succcurlyeq \eta'$: η is "more informative" than η' . $\eta \succcurlyeq \eta'$ if every potential user would prefer η over η' .

It can be shown that [9]:

$$Z \, f \, Z' \Longrightarrow Z \, g \, Z' \Longrightarrow \eta \succcurlyeq \eta' \qquad (2.60)$$

but the converse relations do not hold in general. We will see in the following that some of the relations defined in Definitions (1), (2) and (3) are preserved or reduced to weaker relations as information becomes old.

I. Let us suppose that $Z \, f \, Z'$. Which of the relations (1), (2), or (3) hold for delayed information? Let us denote the delayed information by $y(t) = \xi(s(t))$. We have $y(t) = \xi(s(t)) = \eta(s(t-\tau)) = z(t-\tau)$.

(1) $Y \, g \, Y'$ does not hold (in general) because ξ and ξ' are not noiseless .

(2) $Y \, g \, Y'$ holds because

$$\{y'(t)|y(t),s(t)\mathscr{B}\} = \{z'(t-\tau)|z(t-\tau),s(t),\mathscr{B}\}$$

but by (2.60) $Z\,g\,Z'$ holds and we have:

$$\{z'(t-\tau)|z(t-\tau),s(t),\mathscr{E}\} = \{z'(t-\tau)|z(t-\tau),\mathscr{E}\}$$

Therefore

$$\{y'(t)|y(t),s(t),\mathscr{E}\} = \{z'(t-\tau)|z(t-\tau),\mathscr{E}\}$$

$$= \{y'(t)|y(t),\mathscr{E}\}$$

(3) $\xi \succcurlyeq \xi$ holds by (2.60) because $Y\,g\,Y'$ holds.

II. Let $Z\,g\,Z'$, then for delayed information we have

(1) $Y\,f\,Y'$ clearly does not hold in general.

(2) $Y\,g\,Y'$ holds by the same argument as in I-(2), above.

(3) $\xi \succcurlyeq \xi'$ holds because $Y\,g\,Y'$ holds.

III. Let $\eta \succcurlyeq \eta'$, then for delayed information:

(1) $Y\,f\,Y'$ clearly does not hold in general.

(2) $Y\,g\,Y'$ clearly does not hold in general.

(3) $\xi \succcurlyeq \xi'$ does not hold in general, but it holds for Markovian systems. The following example shows that $\xi \succcurlyeq \xi'$ does not hold in general. Let

$z(t) = \eta(s(t)) = s(t)$ be perfect information

$z'(t) = \eta'(s(t)) = s(t-1)$ be perfect, but one unit delayed information

Clearly $\eta \succ \eta'$, but if both η and η' have one unit delay we have

$$y(t) = \xi(s(t)) = \eta(s(t-1)) = s(t-1)$$

$$y'(t) = \xi'(s(t)) = \eta'(s(t-1)) = s(t-2)$$

and ξ is not, in general, more informative than ξ'. For example, in the autoregressive system of the second order we showed that (Sec. 2.9) if $\alpha_2 > |\alpha_1|$, then knowing $s(t-2)$ is more valuable than knowing $s(t-1)$. To show that $\xi \succ \xi'$ for Markovian systems, let $v(s(t),d)$ be a payoff function, and the Markovian system be represented by

$$s(t) = f(s(t-1))$$

where f is, in general, a one-to-many mapping and can be assumed, without loss of generality, to remain the same over time. We then have

$$s(t-1) = f(s(t-2))$$

$$s(t-2) = f(s(t-3))$$

$$\vdots$$

Therefore,

$$s(t) = f^{(\tau)}(s(t-\tau))$$

where by $f^{(\tau)}$ we mean τ times applying of f. The payoff at time t is

$$v(s(t),d) = v(f^{(\tau)}(s(t-\tau)),d)$$

Now let

$$w(s(t-\tau),d) = \langle v(f^{(\tau)}(s(t-\tau)),d) \,|\, s(t-\tau),\mathscr{E}\rangle$$

Then we have

$$\langle v(s(t),d) \,|\, y(t) = \eta(s(t-\tau)),\mathscr{E}\rangle$$

$$= \int_{s(t-\tau)} \langle v(s(t),d) \,|\, \eta(s(t-\tau)),s(t-\tau)\,|\,,\mathscr{E}\rangle \cdot \{s(t-\tau) \,|\, \eta(s(t-\tau)),\mathscr{E}\}$$

$$= \int_{s(t-\tau)} w(s(t-\tau),d)\{s(t-\tau) \,|\, \eta(s(t-\tau)),\mathscr{E}\}$$

$$= \langle w(s(t-\tau),d) \,|\, \eta(s(t-\tau)),\mathscr{E}\rangle$$

and

$$\max_d \langle v(s(t),d) \,|\, \eta(s(t-\tau)),\mathscr{E}\rangle = \max_d \langle w(s(t-\tau),d) \,|\, \eta(s(t-\tau)),\mathscr{E}\rangle \quad (2.61)$$

and, similarly,

$$\max_d \langle v(s(t),d) \,|\, \eta'(s(t-\tau)),\mathscr{E}\rangle = \max_d \langle w(s(t-\tau),d) \,|\, \eta'(s(t-\tau)),\mathscr{E}\rangle \quad (2.62)$$

but since η is more informative than η', η is more valuable than η' for any payoff function and in particular for $w(s,d)$. It follows that the right-hand side of (2.61) is greater than the right-hand side of (2.62); therefore, the left-hand side of (2.61) is greater than the left-hand side of (2.62):

$$\max_d \langle v(s(t),d) | \xi(s(t)),\mathscr{E} \rangle \geq \max_d \langle v(s(t),d) | \xi'(s(t)),\mathscr{E} \rangle$$

and since $v(s,d)$ is an arbitrary function, it follows that $\xi \succcurlyeq \xi'$.

Our work in this section has been more of an exploratory type. We showed that the relative advantage of one information structure over another is not, in general, preserved when both information structures are delayed. We found, however, that such a relation would be preserved (or somewhat weakened), if the advantage of one information structure over the other is not limited to a particular payoff function.

2.11 Summary

The process of outdating of information in a dynamic environment has been investigated in this chapter. We have shown how this process depends on the various factors which influence it, namely the dynamics of the environment, the decision for which the information is used, and the properties of the information structure itself (perfect or imperfect). Assuming a quadratic payoff function, the value of an information structure was calculated as a function of its age. The result was written in a form which shows separately the effect of each factor on the information outdating process:

$$V_{\eta(\tau)} = \mathrm{tr}\left[M \cdot R(\tau) \sum_{\underset{\sim}{s}_o} \cdot R'(\tau) \right] \qquad (2.63)$$

$V_{\eta(\tau)}$, the value of the information structure η with delay τ , is the trace of the product of four matrices, which represent the factors influencing the outdating of information. The matrix M represents the effect of the payoff function on the value of information (M is determined by the coefficient matrices of the payoff function). The information structure influences the value of information through matrix $\Sigma_{\underset{\sim}{s}_o}$ of the covariance of the posterior mean of the state, given fresh information. Matrices $R(\tau)$ and $R'(\tau)$ (transpose of $R(\tau)$) represent the dynamics of the state (environment). $R(\tau)$ is the coefficient of the linear approximation of the mean of the state at time t , as a function of the state at time $t - \tau$. If the state has a normal distribution, $R(\tau)$ is the "multiple correlation" between $\underline{s}(t)$ and $\underline{s}(t - \tau)$. Since we assume a stationary state, $R(\tau)$ is independent of t .

As we can see from Eq. (2.63), the main determinants of the dynamics of the information are the dynamics of the state, as represented by the matrix $R(\tau)$. Nevertheless, the payoff function and the information structure itself can drastically influence the dynamics of information. Specific results have been obtained for the linear Markovian system,

$$\underline{s}(t) = A \cdot \underline{s}(t-1) + \underline{\varepsilon}(t)$$

where A is a constant matrix and $\underline{\varepsilon}(t)$ is a random noise. We have shown that the eigenvalues of matrix A play an important role in determining the pattern of the information outdating process. It was found that the value of information may increase with delay or may even oscillate. This is a rather counter-intuitive result, especially for a Markovian system. We have shown, however, that if the information is perfect, its value will always decrease with delay, regardless of the dynamics of the state or the parameters of the payoff function. We have found other conditions under which the information is always perishing. For these cases, bounds are found for the rate of information perishing. These bounds are determined by the smallest and the largest eigenvalues of matrix A. The results for the Markovian system are summarized in Table 2.1.

Table 2.1 Dynamics of Information in a Linear Markovian System

Properties of the state, the information, and the decision	Dynamics of the Information
η Perfect	Information always perishing
$A, M, \Sigma_{\underset{\sim}{s}_o} \geq 0$	Information always perishing; $\rho(\tau) \geq \frac{1}{\lambda_N^2}$ (a)
A and M diagonal	Information always perishing; $\frac{1}{\lambda_1^2} \geq \rho(\tau) \geq \frac{1}{\lambda_N^2}$ (a)
A and $\Sigma_{\underset{\sim}{s}_o}$ diagonal	Information always perishing; $\frac{1}{\lambda_1^2} \geq \rho(\tau) \geq \frac{1}{\lambda_N^2}$ (a)
A has complex eigenvalues	The value of the information may oscillate with delay.

(a) λ_1 and λ_N are the smallest and the largest eigenvalues of matrix A.

We have also investigated the information outdating process for the autoregressive systems. Although an autoregressive system can be represented by a Markovian system of a higher dimension, the study was useful in gaining insight into the process of information outdating. In particular, it helped to show more clearly why the value of the information may be enhanced or may oscillate with time.

Finally, we have made some comparisons between delayed information structures. We have investigated whether various relations between two information structures (regarding the superiority of one over another) are preserved, when both information structures are delayed. The relations studied were the following: (1) $Z f Z'$: set Z of observation η is finer than set Z' of observation η' ; (2) $Z g Z'$: set Z of observation η is "garbled" into set Z' of observation η' ; (3) $\eta \succcurlyeq \eta'$: observation η is "more informative" than observation η' . The first relation is the strongest (with regards to η being superior to η'), and the third relation is the weakest. When the information structures η and η' are delayed some of the relations are preserved, and some are reduced to weaker relations. The results are summarized in Table 2.2.

Table 2.2. Comparison Between Delayed Information Structures.

Fresh Information	Delayed Information
$Z_o \ f \ Z'_o$	$Z_\tau \ f \ Z'_\tau$ does not hold $Z_\tau \ g \ Z'_\tau$ holds $\eta_\tau \geqslant \eta'_\tau$ holds
$Z_o \ g \ Z'_o$	$Z_\tau \ g \ Z'_\tau$ holds $\eta_\tau \geqslant \eta'_\tau$ holds
$\eta_o \geqslant \eta'_o$	$\eta_\tau \geqslant \eta_\tau$ does not hold in general but holds for Markovian systems.

CHAPTER 3

RANDOMLY OCCURRING DECISIONS AND INFORMATION RECOVERY

In Chapter 2 the process of outdating of information was studied. There are many decision situations where we must deal with old (unfresh) information. In such cases, obtaining fresh information at the time of the decision may be very time consuming or extremely expensive. One of the most important cases, where the use of unfresh information may be inevitable, is when the time of the decision is uncertain. The decision may or may not have to be made at each point of time. Moreover, when the time of the decision comes, it must be made within a short period of time, thus making it impossible to obtain fresh information for the decision. We may think of this type of decision as one which must be made upon the occurrence of a precipitating event with random time. Often the decision maker has no control over this event. Following Grum [5] we use the term "contingent" for this type of decision. The precipitating event in a contingent decision is often an action by another party. For example, firms have to make contingent decisions in response to either government actions (e.g. regulations, bidding contracts, etc.), or actions by their competitors (e.g. changing prices, introducing new products, etc.). The government is faced with contingent decisions, often as a result of economic or political decisions by foreign countries (e.g. outbreak of wars, economic embargos, etc.). Since we are dealing with unfresh information in making such decisions, and in view of the information being outdated in time, it may be desirable to be prepared for the decision by regular recovery (updating) of information.

In this chapter two approaches to the recovery of information are introduced. The optimal recovery of information according to these approaches is studied in Chapters 4 and 5.

3.1 The Contingency Decision Model

As mentioned before, a contingent decision may be thought of as a decision which must be made upon the occurrence of a precipitating event with random time. We denote this event by E. E may be a single event, a combination of simultaneous events, or the last of a chain of events. In any case, a single event E can represent all these situations. In Chapters 4 and 5 we study the one-time decision case, namely the case where the decision happens only once. In Chapter 6 the results are extended to the case where the decision may be repeated in time. For the one-time decision case the occurrence of the decision may be modeled as shown in Fig. 3.1. g_t is the probability that event E occurs at time t (given that it did not occur before). Of course, g_t may be revised as time passes. E occurs only once.

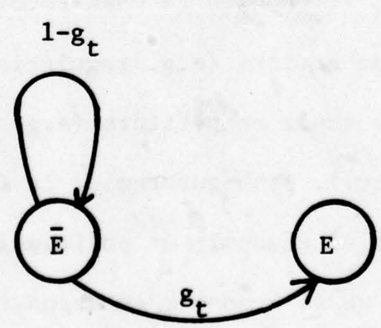

Figure 3.1 Occurrence model for a contingent decision.

3.2 Information Recovery: A priori vs. A posteriori Policies

As mentioned earlier, for contingent decisions we may find it attractive to update our information regularly in order to be prepared for the decision. Two types of policies for information recovery (updating) are studied here.

1. A priori policies for information recovery: In this type of policy we use our prior knowledge about the outdating of information to decide the times of all future observations (recoveries). In particular, if g_t (the probability of the decision occurring at time t given that it did not occur before) is constant and our system is stationary, the information recovery for the infinite-horizon case will be periodic, as shown in Fig. 3.2. In this figure the information is recovered

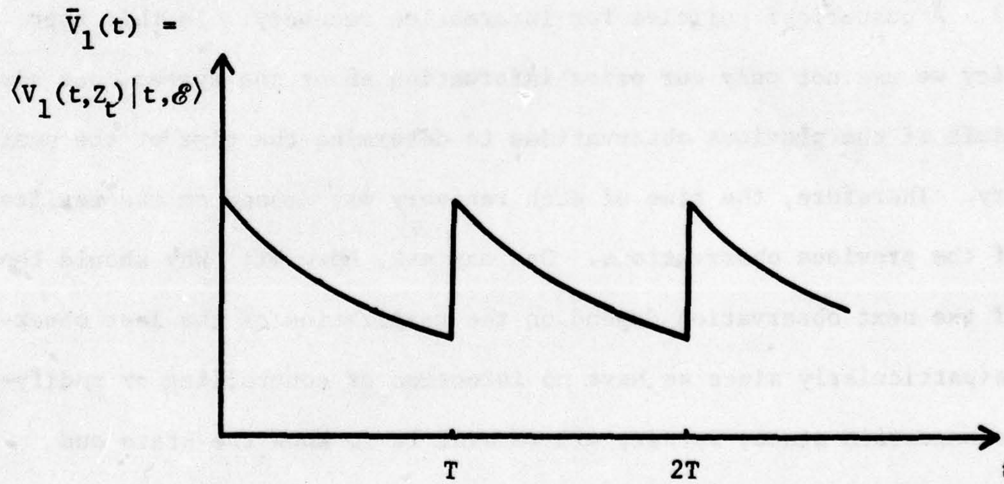

Figure 3.2 Apriori Recovery of Information.

at times $0, T, 2T, \ldots$. $\bar{V}_1(t)$ is the a priori expected payoff if the decision occurs at time t. Note that the actual expected payoff at time t depends

on the information available at time t , namely the result of the observations before t . Denoting the sequence of all the observations before t by Z_t , we have

$$\bar{V}_1(t) = \langle V_1(t,Z_t) | t,\mathscr{E} \rangle$$

where $V_1(t,Z_t)$ is the expected payoff at time t , given Z_t . We can assume, without loss of generality, that the expected payoff with no information is zero. Thus, $\bar{V}_1(t)$ will be the a priori value of the past information at time t . Therefore, this type of policy concerns the a priori outdating of information and appropriate information recovery schedules.

2. A posteriori policies for information recovery: In this type of policy we use not only our prior information about the system, but also the result of the previous observations to determine the time of the next recovery. Therefore, the time of each recovery may depend on the realization of the previous observations. One may ask, however: Why should the time of the next observation depend on the realization of the last observations (particularly since we have no intention of controlling or modifying the uncertain state; rather, all we want is to know the state and then set our decision accordingly)? The following example illustrates this matter. Consider a target in a field at which we like to shoot if a random event E occurs (assume E to be independent of the position or any other property of the target). The target is constantly moving in an uncertain manner and we cannot see its location. We can find out

about its location at any time, however, at some cost. When E occurs, we have to shoot at the target immediately (using only our previously obtained information about its location), and our payoff depends on how closely we hit it. If all that matters is how closely we hit the target regardless of where the target is in the field, and if the motion of the target is "independent" of its position, then there is no reason why the time of the next observation of the location of the target should depend on its location in the previous observations. All we want is to know where the target is, but its position per se is of no significance. However, if the field is not "homogeneous" in the sense that different regions in the field have different degrees of importance, then the time of the next observation may well depend on the locations observed previously. Suppose, for example, that one region of the field is of particular importance and therefore there is a very high payoff for hitting the target in that region. Then if the last observation shows that the target is in this region we may want to observe its location more frequently. We will elaborate more on this in Chapter 5, where we find some conditions which make the time of the next recovery independent of the realization of the last observations.

The a posteriori information recovery problem may be thought of as deciding at each period whether or not to buy information at that period (given the state of information at that period). This is shown in Fig. 3.3. At each time t^- (recall that t is discrete, t^- denotes slightly before t) we must decide whether or not to buy new information,

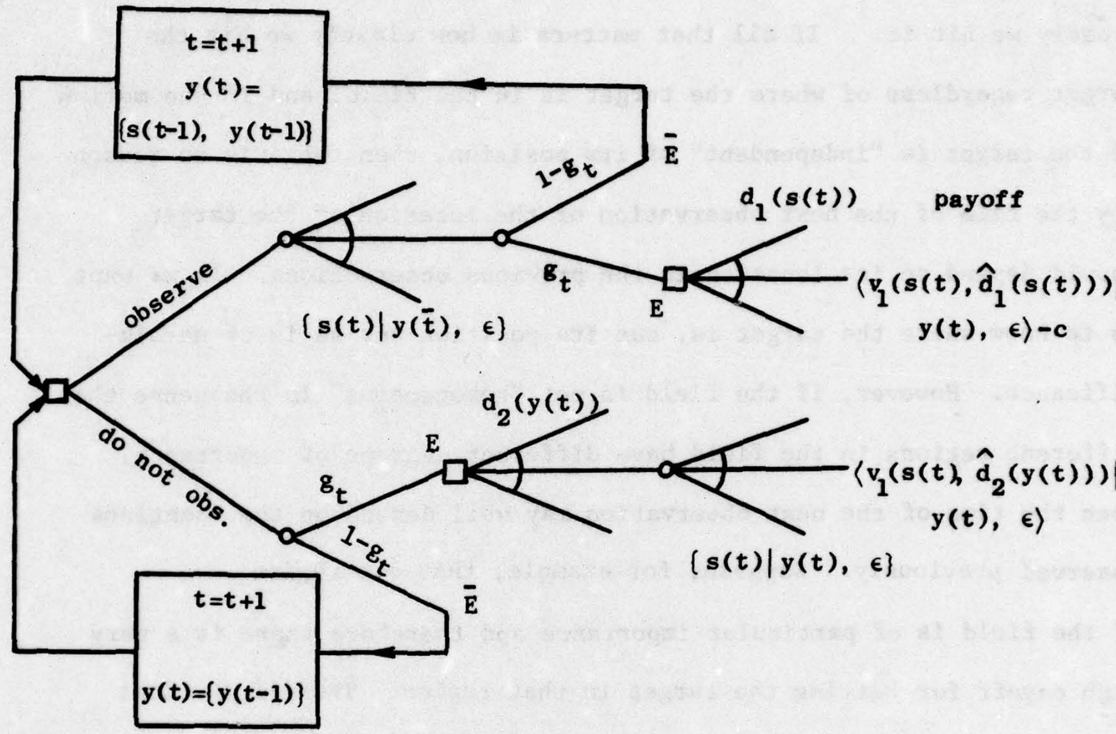

Figure 3.3 Decision model for the a posteriori recovery of information.

given our state of information at t^- $(y(t^-))$. If we buy information we will learn $s(t)$ (assuming perfect information) with the distribution $\{s(t)|y(t^-), \mathscr{E}\}$. If the event E happens at t , then we will set our decision d_1 according to the new information $s(t)$, and the net payoff will be $\langle v_1(s(t),\hat{d}_1(s(t))|y(t^-),\mathscr{E}\rangle - c$, where $\hat{d}_1(s(t))$ denotes the optimal decision given $s(t)$, and c is the cost of information.

If E does not happen, we will proceed to time $t+1$ where we are faced with the same decision, but with a new state of information updated at time t. If we do not acquire information at t^-, then if E happens we have to set our decision d_2 according to our previous information $y(t)$ and the payoff will be $\langle v_1(s(t),\hat{d}_2(y(t))|y(t),\mathcal{E}\rangle$. If E does not happen, we proceed to time $t+1$ where we are faced with the same decision again, but our information is one unit older.

Note that since no new information is obtained during the intervals between observations, we may try to calculate the next recovery time immediately after each observation. Therefore, we may also think of our information recovery problem as deciding the next recovery time after observing the state at each recovery.

For a Markovian state with perfect observations, the result of the last observation of the state is all the information needed for deciding the next recovery time. Consequently, the calculations of the optimal policy will be greatly reduced.

The a posteriori policies clearly have a higher expected payoff compared to the a priori policies. They are more difficult to calculate, and perhaps more costly to implement, however.

In Chapter 4 we will investigate the optimal a priori policies for information recovery. The a posteriori policies are studied in Chapter 5.

CHAPTER 4

A PRIORI OPTIMAL INFORMATION RECOVERY

In this chapter we will study the a priori optimal information recovery policies for contingent decisions. The optimality conditions are found for finite and infinite horizon cases. The effect of risk aversion on the optimal information recovery policies is also investigated.

4.1 Optimal Information Recovery for the Infinite-Horizon Case

Most of our results concern the infinite horizon case. Assuming that $s(t)$ is stationary and g_t (the probability that the precipitating event E occurs at time t, given that it did not occur before) is constant, the optimal information recovery is periodic. We denote the recovery period by T. $\bar{V}_1(t)$ was defined in Chapter 3 as the a priori expected payoff of the decision given that it occurs at time t. We have

$$\bar{V}_1(t) = \langle V_1(t, Z_t) | t, T, \mathcal{E} \rangle \qquad (4.1)$$

where $V_1(t,Z_t)$ is the posterior expected payoff if the decision occurs at t (Z_t is the sequence of all the observations before t). Since $\bar{V}_1(t)$ is periodic (Fig. 4.1), it is easier to measure t from a recovery time. Therefore we will consider $\bar{V}_1(t)$ as the a priori expected payoff t time units after a recovery.

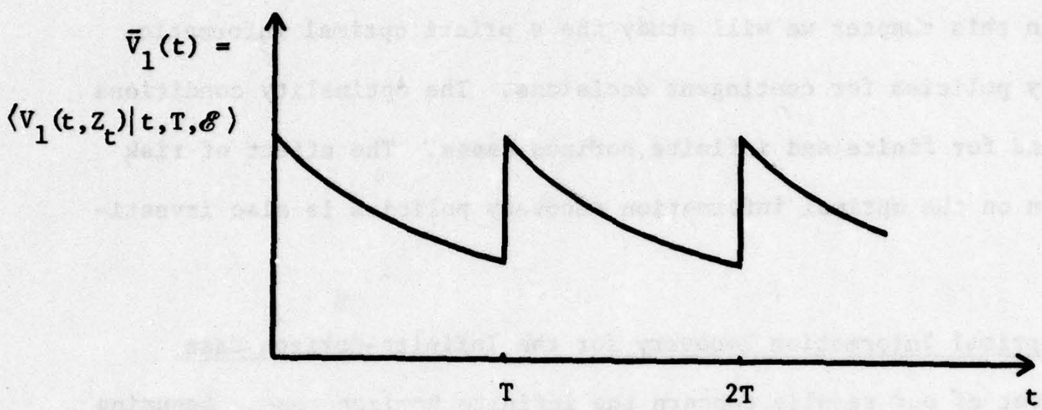

Figure 4.1 Information recovery for the infinite-horizon case.

In a contingent decision the time of the decision is uncertain. We define $\bar{V}(T)$ as the <u>net</u> expected payoff given only that the recovery period is T (but with the time of the decision uncertain). We have

$$\bar{V}(T) = \langle V(t,Z_t)|T,\mathscr{E}\rangle \tag{4.2}$$

where $V(t,Z_t)$ is the net expected payoff if the decision occurs at time t, and Z_t was observed in the previous observations. Notice that since $V(t,Z_t)$ is the net payoff it takes into account the cost of information as well as the payoff of the decision.

<u>Theorem 4.1.</u> If

(1) Information is always perishing (that is $\bar{V}_1(t)$ is a decreasing function of t), and

(2) There exists at least one T, such that the net expected payoff with recovery period T ($\bar{V}(T)$) is positive. Then, there exists a unique optimum recovery period T^* (See Fig. 4.2).

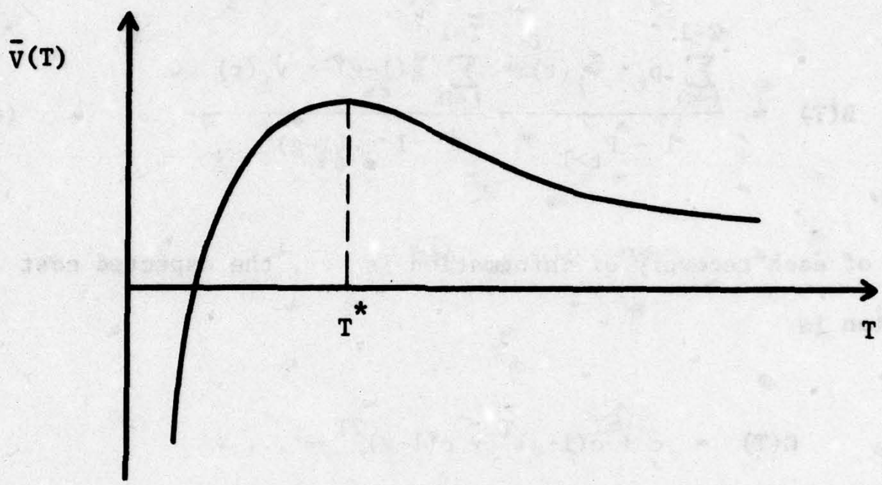

Figure 4.2 Expected payoff as a function of the recovery period.

<u>Proof</u>: Since the horizon is infinite, the expected payoff (benefit) with recovery period T, $B(T)$, can be written as

$$B(T) = \sum_{t=0}^{T-1} p_t \cdot \bar{V}_1(t) + P_{t \geq T} \cdot B(T) \qquad (4.3)$$

where $p_t = g(1-g)^t$ is the probability (at time zero) that the decision will occur at time t, and $P_{t \geq T}$ is the probability that the decision will not occur before T. The first term on the right-hand side of (4.3) is the expected payoff in the first period and the second term is the

expected payoff after the first period. Note that we have not discounted future payoffs. This would only be done for simplicity, and there is no difficulty in using the discounted values. From (4.3) we have

$$B(T) = \frac{\sum_{t=0}^{T-1} p_t \cdot \bar{V}_1(t)}{1 - P_{t \geq T}} = \frac{\sum_{t=0}^{T-1} g(1-g)^t \cdot \bar{V}_1(t)}{1 - (1-g)^T} \tag{4.4}$$

If the cost of each recovery of information is c , the expected cost of information is

$$\begin{aligned} C(T) &= c + c(1-g)^T + c(1-g)^{2T} + \ldots \\ &= \frac{c}{1 - (1-g)^T} \end{aligned} \tag{4.5}$$

The net expected payoff is therefore

$$\bar{V}(T) = B(T) - C(T) = \frac{\sum_{t=0}^{T-1} g(1-g)^t \ \bar{V}_1(t) - c}{1 - (1-g)^T} \tag{4.6}$$

Now we can show that

(1) If $\bar{V}(T_o) > \bar{V}(T_o+1)$, then $\bar{V}(T_o+1) > \bar{V}(T_o+2)$, and

(2) If $\bar{V}(T_o) < \bar{V}(T_o+1)$, then $\bar{V}(T_o-1) < \bar{V}(T_o)$.

The first statement implies that if $\bar{V}(T)$ is decreasing from T_o to T_o+1 , it will be decreasing for all $T > T_o$. The second statement

implies that if $\bar{V}(T)$ is increasing from T_o to T_o+1, it is increasing for all $T < T_o$. From (1) and (2) it follows that if $\bar{V}(T)$ has a maximum, it will be unique. To prove (1) and (2), let us find the expression for $\bar{V}(T) - \bar{V}(T+1)$. Letting $\bar{V}_1(t) = \alpha_t \cdot \bar{V}_1(0)$ we have

$$\bar{V}(T) - \bar{V}(T+1) = \frac{g(1-g)^T}{[1-(1-g)^T][1-(1-g)^{T+1}]} \cdot G(T) \quad (4.7)$$

where

$$G(T) = g\bar{V}_1(0)\left[1 + \alpha_1(1-g) + \alpha_2(1-g)^2 + \ldots + \alpha_{T-1}(1-g)^{T-1}\right]$$
$$- \alpha_T \bar{V}_1(0)\left[1-(1-g)^T\right] - c \quad (4.8)$$

Since the term $g(1-g)^T/\{[1-(1-g)^T][1-(1-g)^{T+1}]\}$ is positive, it is sufficient to show that

(1) $G(T) > 0 \Longrightarrow G(T+1) > 0$

(2) $G(T) < 0 \Longrightarrow G(T-1) < 0$

$G(T+1)$ can be written as

$$G(T+1) = G(T) + g\,\bar{V}_1(0)(\alpha_T^2 - \alpha_{T+1}^2)\left[1-(1-g)^{T+1}\right]$$

but by the information perishing assumption $\alpha_T > \alpha_{T+1}$, and the second term on the right-hand side is positive; therefore,

$$G(T) > 0 \Longrightarrow G(T+1) > 0$$

Similarly, we have

$$G(T-1) = G(T) + g\,\bar{V}_1(0)(\alpha_T^2 - \alpha_{T-1}^2)[1 - (1-g)^T]$$

and since $\alpha_T < \alpha_{T-1}$, the second term on the right-hand side is negative, and we have

$$G(T) < 0 \implies G(T-1) < 0 \ .$$

Now note that $G(T)$ can be written as

$$G(T) = (1 - (1-g)^T)[\bar{V}(T) - \alpha_T \bar{V}_1(0)] \tag{4.9}$$

as $T \to \infty$, $\alpha_T \bar{V}_1(0) \to 0$ and $G(T)$ will have the same sign as $\bar{V}(T)$. Therefore, if $\bar{V}(\infty) > 0$, then $G(\infty) > 0$ and $\bar{V}(T)$ is decreasing for large T . Similarly $V(\infty) \leq 0$ implies $\bar{V}(T)$ increasing for large T . From this observation and from (1) and (2) it follows that $\bar{V}(T)$ has one of the forms of Fig. 4.3 (noting that $V(0) = -\infty$) . This completes the proof of the theorem.

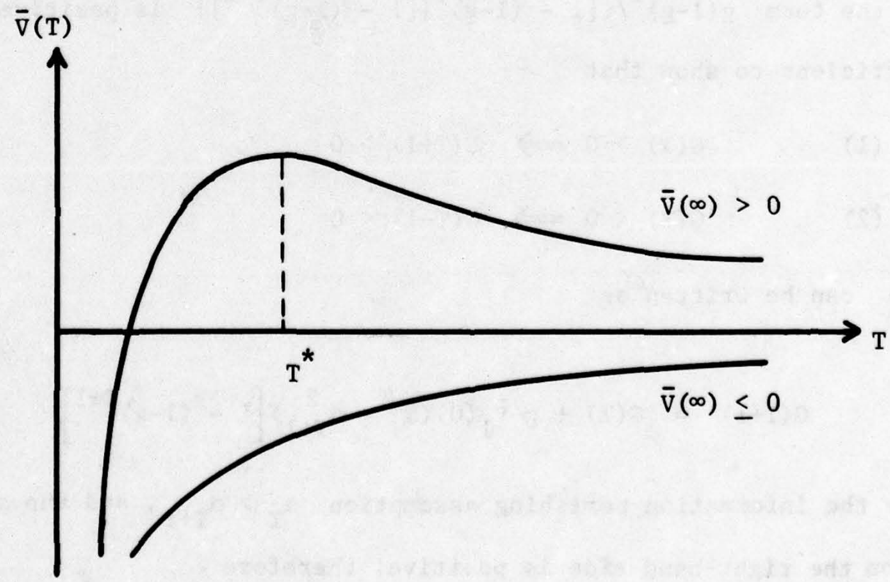

Figure 4.3 Optimal recovery period T^* .

__Theorem 4.2__. The necessary and sufficient condition for the optimal information recovery period (T^*) is that the residual value of information immediately before information recovery $(\bar{V}_1(T^*))$ be equal to the expected net payoff of the policy $\bar{V}(T^*)$.

__*Proof:*__ We showed in Theorem 4.1 that $G(T)$, as defined in (4.8), is negative if $\bar{V}(T)$ is increasing and is positive if $V(T)$ is decreasing. From Fig. 4.3 it follows that $G(T)$ has one of the forms shown in

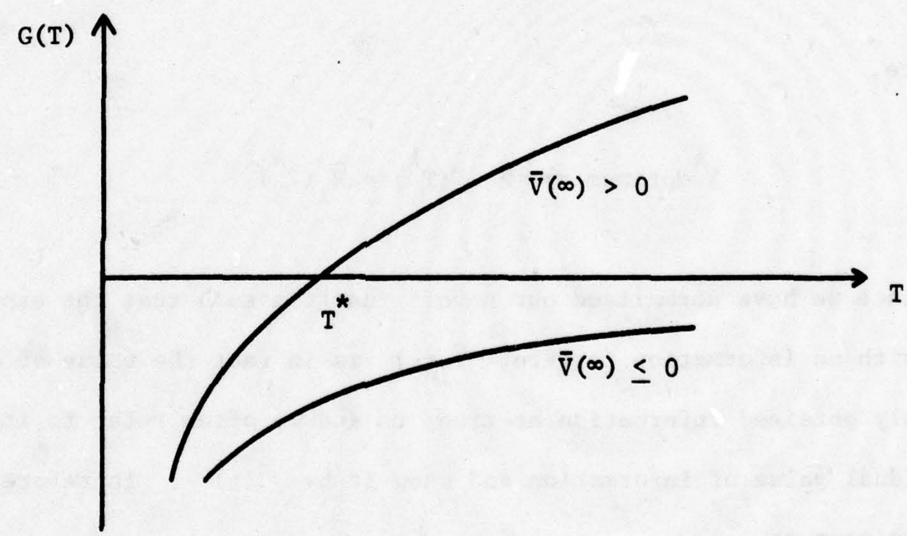

Figure 4.4 $G(T)$ for $\bar{V}(\infty) > 0$ and $\bar{V}(\infty) \leq 0$.

Fig. 4.4. If $\bar{V}(\infty) > 0$, then T^* is unique and finite and we must have:

$$G(T^*) = 0 \tag{4.10}$$

This is the necessary and sufficient condition for T^* because of the uniqueness of T^*. If $\bar{V}(\infty) < 0$, there is no optimum and $G(T)$

is always negative. For $\bar{V}(\infty) = 0$, $G(T)$ is zero at $T = \infty$. $T^* = \infty$ may be considered as optimum although the payoff is the same as buying no information at all. Therefore (4.10) is always a necessary and sufficient condition for the optimal information recovery period T^* . From (4.9) and (4.10) we have

$$\bar{V}(T^*) = \alpha_{T^*} \cdot \bar{V}_1(0) = \bar{V}_1(T^*) \tag{4.11}$$

Therefore,

$$T^* \text{optimum} \Longleftrightarrow \bar{V}(T^*) = \bar{V}_1(T^*)^{\dagger} \tag{4.12}$$

Since we have normalized our payoff function such that the expected payoff with no information is zero, $\bar{V}_1(t)$ is in fact the value of our previously obtained information at time t, and we often refer to it as the residual value of information and show it by $\bar{R}(t)$. Therefore (4.12) may be written as

$$\bar{V}(T^*) = \bar{R}(T^*) \tag{4.13}$$

According to Theorem 4.2 the net expected payoff of a contingency decision is at most equal to the (gross) expected payoff if the decision

---

† Note that T^* which satisfies Eq. (4.11) is not, in general, an integer. The time in our system is discrete, however. Since $G(T) = 0$ implies $\bar{V}(T) = \bar{V}(T+1)$, it follows that the actual optimum T_* (integer) lies between T^* (from Eq. (4.11)) and T^*+1 . Therefore T^* found in (4.11) must be rounded up to find the actual optimum T , but we often ignore the discrepancy and regard T^* as optimum.

occurs when our information is at its lowest point (Fig. 4.5). The extra payoff of the decision if it occurs at other points of time will just be enough to compensate for the cost of information. We can also make the

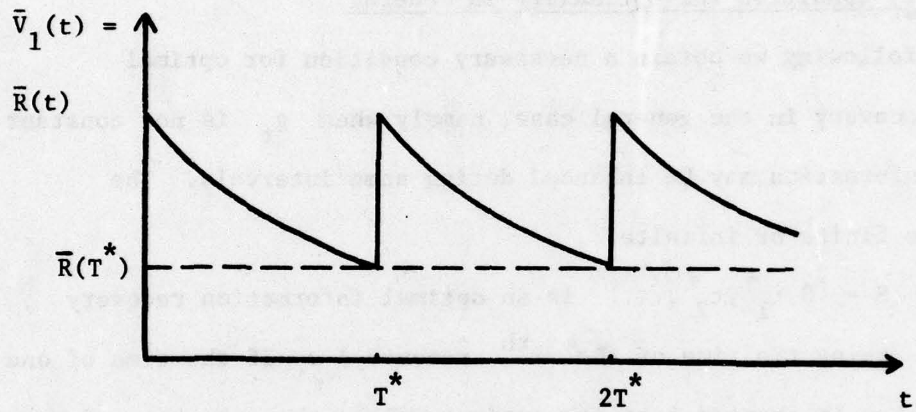

Figure 4.5 $\bar{R}(T^*)$, the residual value of information at T^* .

following remarks from Theorem 4.2.

Remark 1. The net expected payoff of the contingency decision with an optimal recovery period (T^*) is the same as the expected payoff in a normal decision (known time) with information having a delay T^* .

This is immediate from (4.12) since $\bar{V}_1(t)$ is the expected payoff with the information having a delay t .

Remark 2. If the optimal information recovery period T^* increases (as a result, for instance, of an increase in the cost of information), then the net expected payoff (with optimum recovery period) decreases at the same rate as the rate of information perishing of the system. This is again immediate from (4.12).

Remark 3. If there is any information recovery it will be before the past information is completely perished.

From (4.12), $\bar{V}(T^*) > 0$ if and only if $\bar{V}_1(T^*) > 0$, that is, if the information is not completely perished at recovery times.

4.2 Necessary Condition for Optimality in General

In the following we obtain a necessary condition for optimal information recovery in the general case, namely when g_t is not constant in time and information may be enhanced during some intervals. The horizon may be finite or infinite.

Suppose $S = \{0, t_1^*, t_2^*, \ldots\}$ is an *optimal information recovery schedule* (t_i being the time of the i^{th} recovery). If the time of one of the recoveries is changed from its optimum value, the net expected payoff decreases (or remains unchanged), no matter how the timing of the other recoveries are modified. Let us define

$$\bar{V}_t(t_i) = \text{maximum net expected payoff from time } t \text{ on with the next recovery at } t_i (t_i \geq t) \text{, given the previous recoveries schedule.} \tag{4.14}$$

$\bar{V}_t(t_i)$ reaches its maximum at $t_i = t_i^*$ for all $t \in [t_{i-1}^*, t_i^*]$ (See Fig. 4.6). This is true because our state of information does not change in the interval (t_{i-1}^*, t_i^*). Now note from Fig. 4.6 that the equation

$$\bar{V}_t(\hat{t}_i) = \bar{V}_t(\hat{t}_i + 1) \tag{4.15}$$

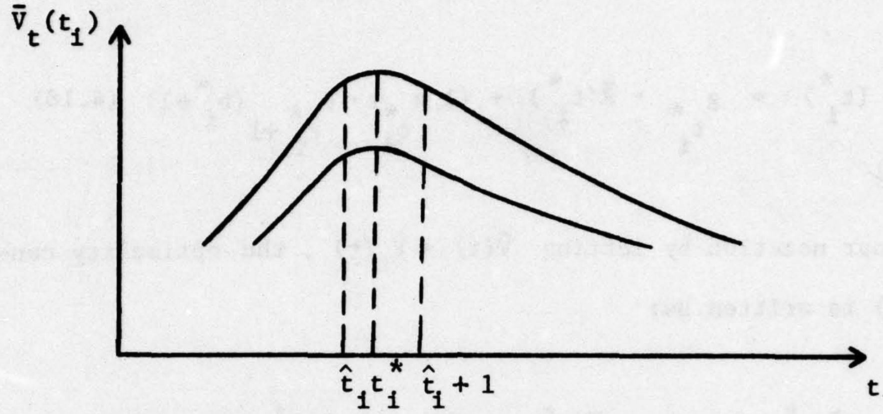

Figure 4.6 Optimum information recovery time t_i^* .

can hold only for some $\hat{t}_i$, such that $\hat{t}_i < t_i^* < \hat{t}_i + 1$. Since t_i^* must be an integer (because time is discrete in our system), t_i^* is in fact equal to $\hat{t}_i$ when the latter is rounded up. Therefore (4.15) can be regarded as the equation for optimum t_i^* , keeping in mind that the t_i^* obtained from this equation must be rounded up. Therefore we write

$$\bar{V}_t(t_i^*) = \bar{V}_t(t_i^* + 1) \qquad (4.16)$$

Letting $t = t_i^*$, $\bar{V}_t(t_i^* + 1)$ can be written as

$$\bar{V}_{t_i^*}(t_i^* + 1) = g_{t_i^*} \cdot \bar{R}(t_i^*) + (1 - g_{t_i^*}) \cdot \bar{V}_{t_i^*+1}(t_i^* + 1) \qquad (4.17)$$

$\bar{R}(t)$ is the residual value of information at time t ; therefore, the first term on the right-hand side of (4.17) is the expected payoff if the decision occurs at t_i^* . The second term is the expected payoff from

time $t_i^* + 1$ on. From (4.16) and (4.17) we have

$$\bar{V}_{t_i^*}(t_i^*) = g_{t_i^*} \cdot \bar{R}(t_i^*) + (1-g_{t_i^*}) \cdot V_{t_i^*+1}(t_i^*+1) \quad (4.18)$$

Simplifying our notation by letting $\bar{V}(t) = \bar{V}_t(t)$, the optimality condition (4.18) is written as:

$$\bar{V}(t^*) = g_{t^*} \cdot \bar{R}(t^*) + (1-g_{t^*}) \cdot \bar{V}(t^*+1) \quad (4.19)$$

Recall that

$\bar{V}(t)$ = maximum expected payoff from t on with recovery at t .

and

$\bar{R}(t)$ = residual value of information at t .

Notice also that for the special case of the previous section, namely when g_t is constant and the horizon is infinite, $\bar{V}(t^*) = V(t^*+1)$, and (4.19) reduces to

$$\bar{V}(t^*) = \bar{R}(t^*)$$

which is the condition we had obtained for that case.

4.3 Interpretation of the Optimality Condition

$\bar{V}(t)$ can be written as

$$\bar{V}(t) = g_t \cdot \bar{V}_1'(t) + (1-g_t)[\bar{V}_\delta(t+1|t) + c] - c \quad (4.20)$$

where $\bar{V}_1'(t)$ is the expected payoff if the decision occurs at t with information being recovered at t, and $\bar{V}_\delta(t+1|t)$ is the net expected payoff from $t+1$ on, given that information is recovered at t (and the decision did not occur at t). c is the cost of information recovery. Substituting for $\bar{V}(t^*)$ from (4.19) into (4.20) we obtain

$$g_{t^*} \cdot [\bar{V}_1'(t^*) - \bar{R}(t^*) - c] = (1-g_{t^*}) \cdot [\bar{V}(t^*+1) - \bar{V}_\delta(t^*+1|t^*)] \qquad (4.21)$$

The left-hand side of (4.21) can be thought of as the net expected loss at t^* of buying information slightly later (t^*+1 rather than t^*) The right-hand side of (4.21) is the net expected benefit from t^*+1 on, of buying information at t^*+1 rather than at t^*. Therefore, the optimality condition (4.21) states that t^* is such that the expected marginal loss at present (t^*) of buying information slightly later (t^*+1 rather than t^*) is equal to its expected marginal benefit in the future (from t^*+1 on).

If information is always perishing, the right-hand side of (4.21) is always positive. Therefore

$$\bar{V}'_1(t^*) - \bar{R}(t^*) - c \geq 0$$

or

$$V_1'(t^*) - R(t^*) \geq c \qquad (4.22)$$

This inequality states that the net benefit at t^* of buying information at t^* (assuming that the decision occurs at t^*) is always greater than the cost of information. Intuitively, if the purchase of information

at time t will not provide a net benefit for making the decision at time t, there is no value in buying it at t, since we can always buy it at $t+1$ without being worse off (assuming perishing information). This argument is valid if the decision occurs only once. When a decision is repeated in time, as we will see in Chapter 5, each purchase of information may be used for more than one decision and (4.22) is not true anymore.

4.4 A Priori Information Recovery in the Bidding Example

The perishing of information in the Bidding Example was studied in Chapter 2 (Sec. 2.8). We found that the information about the cost of performing the bidding contract (p) perishes at a constant rate Eq. (2.49):

$$\bar{V}_1(t) = \frac{\sigma_s}{8} \cdot \lambda^{2t} = \bar{V}_1(0)\, \lambda^{2t} \tag{4.23}$$

λ is a constant less than one. Suppose that the bidding occurs only once and the probability that it occurs at any time t (given that it did not occur before) is constant (g). Assuming an infinite horizon we can use the optimality condition (4.12) to find the optimal information recovery period. We have

$$\bar{V}(T^*) = \bar{V}_1(T^*)$$

$\bar{V}(T)$, namely the net expected payoff with recovery period T from (4.6) is

$$\bar{V}(T) = \frac{\sum_{t=0}^{T-1} g(1-g)^t \cdot \bar{V}_1(t) - c}{1 - (1-g)^T}$$

Substituting for $\bar{V}_1(t)$ from (4.23) we have

$$\bar{V}(T) = \frac{g\, \bar{V}_1(0) \sum_{t=0}^{T-1} (1-g)^t \lambda^{2t} - c}{1 - (1-g)^T}$$

$$= \frac{g\bar{V}_1(0)[1 - ((1-g)\lambda^2)^T] - c}{[1 - (1-g)^T][1 - (1-g)\lambda^2]}$$

Substituting $\bar{V}(T)$ into the optimality condition (4.12) and simplifying, we find

$$\lambda^{2T^*}[1 - (1-u)(1-g)^{T^*}] = \frac{u\,\bar{V}_1(0) - c}{\bar{V}_1(0)} \qquad (4.24)$$

where $u = g/[1 - (1-g)\lambda^2]$ is constant $(0 \leq u \leq 1)$. Since $(1-u)(1-g)^{T^*}$ is often small compared to 1, we have

$$\lambda^{2T^*} \cong \frac{u\,\bar{V}_1(0) - c}{\bar{V}_1(0)}$$

or

$$T^* \cong \frac{\log \left[\frac{u\bar{V}_1(0) - c}{\bar{V}_1(0)}\right]}{\log \lambda^2}$$

but since the rate of information perishing (ρ) is $1/\lambda^2$, we have

$$T^* \cong \frac{\log \dfrac{u\,\bar{V}_1(0) - c}{\bar{V}_1(0)}}{\log \dfrac{1}{\rho}} \tag{4.25}$$

It is easy to show that $u\,\bar{V}(0) - c$ is the net expected payoff if the information is bought only once. Therefore T^* is determined from (1) the value of information when bought only once (normalized by the value of fresh information $\bar{V}_1(0)$), and (2) the rate of information perishing ρ. We buy more information (T^* smaller) if the value of information (when bought only once) increases or if the rate of information perishing (ρ) increases. An increase in the cost of information (c) is reflected directly in the value of information (when bought only once) and increases T^*. The changes in T^* as a result of changes in g and λ (for constant c) can be seen in Fig. 4.7. An increase in g would always result in buying more information, since the value of information (when bought only once) increases. Changes in the value of λ will influence T^* in two opposite ways. Suppose λ decreases, this results in an increase in the rate of information perishing which tends to decrease T^*. On the other hand, the decrease in λ would lower the value of information (when bought only once) which tends to increase T^*. For large values of g the first effect is dominant. The second effect becomes dominant, however, when g is small (Fig. 4.7).

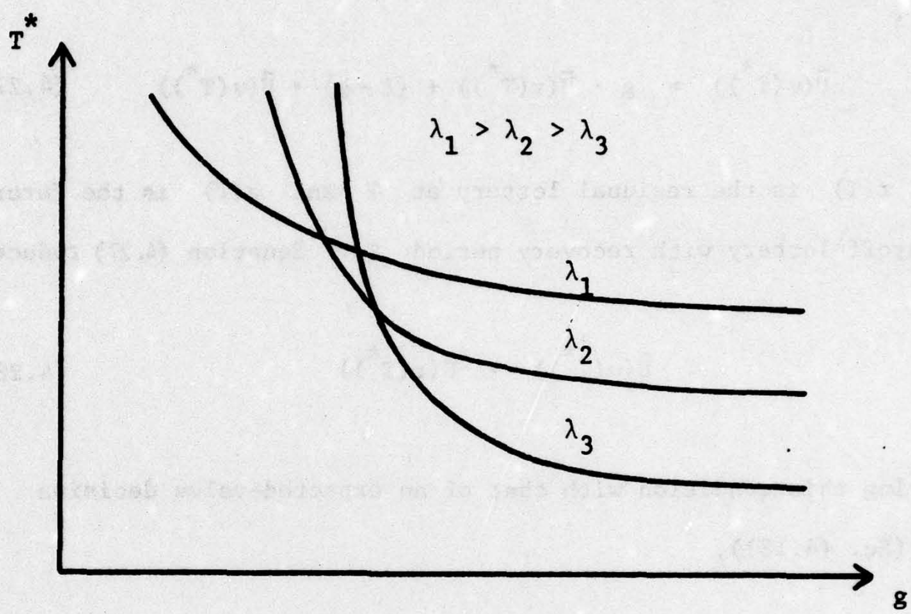

Figure 4.7. T^* as a function of g for various values of λ .

4.5 Optimal Information Recovery with Risk Aversion

The optimality condition for information recovery found earlier in this chapter for an expected value decision maker can easily be extended to the case of a risk-averse decision maker. By the same argument as in Section 4.2, but with expected utilities substituting for expected values, the optimality condition (4.16) is written as

$$\bar{U}_t(t_i^*) = \bar{U}_t(t_i^* + 1) \qquad (t \le t_i^*) \qquad (4.26)$$

where $\bar{U}_t(t_i)$ is the expected utility of the payoff from time t on, with the next recovery at t_i . The optimality condition

for the infinite-horizon case can be obtained by exactly the same method as in the case of an expected-value decision maker. We find

$$\bar{U}(v(T^*)) = g \cdot \bar{U}(r(T^*)) + (1-g) \cdot \bar{U}(v(T^*)) \qquad (4.27)$$

where $r(T)$ is the residual lottery at T and $v(T)$ is the future net payoff lottery with recovery period T. Equation (4.27) reduces to

$$\bar{U}(u(T^*)) = \bar{U}(r(T^*)) \qquad (4.28)$$

Comparing this condition with that of an expected-value decision maker (Eq. (4.13)),

$$\bar{V}(T^*) = \bar{R}(T^*)$$

We note that for an expected-value decision maker, the expected values of the residual lottery and the future payoff lottery are equal (at T^*), whereas for a risk-averse decision maker the utilities of the two lotteries must be equal. Note, however, that the two decision makers are not faced with the same lotteries. In other words, both $r(T)$ and $v(T)$ are different for the two decision makers (because their optimum decisions $(\hat{d})$ are different). This makes any comparison of the optimal information recovery policies between the two cases very difficult. In particular, there is no easy way to discover whether (or under what conditions) a risk-averse decision maker will buy

more or less information than an expected-value decision maker. For a simple Markovian state and a quadratic payoff function, the optimal information recovery periods for the two decision makers were calculated and compared for various values of the parameters involved. In all the cases, where the calculation was done, the information recovery period for the risk-averse decision maker was found to be larger than the one for the expected-value decision maker. In other words, a risk-averse decision maker buys less information than an expected-value decision maker, when there is uncertainty regarding the time of the decision. Whether (or under what conditions) this is in fact true remains uncertain.

4.6 Summary

In this chapter we have investigated the a priori optimal information recovery policies for a one-time contingent decision. Most of our results concern the infinite-horizon case. For this case it was shown that an optimal information recovery period obtained from local optimality conditions would be a global optimum, if information is always perishing. This condition is generally required to make a local optimum a global one. A necessary condition for optimality was found for an information recovery policy in general. This condition can be interpreted as follows: At an optimal information recovery time the expected marginal loss at present (the recovery time) of buying information slightly later is equal to its expected marginal gain in the future. For the infinite horizon case the optimality condition reduces to a simple form and states that the net expected payoff of the decision is equal to the expected payoff of the decision, if it would occur when the information is at its lowest point

(immediately before each recovery). The extra payoff of the decision, if it occurs at other points of time, will be just enough to pay for the cost of the information.

The optimal information recovery period for the Bidding Example of Chapter 2 was found to be determined by (1) the value of information, when it is bought only once, and (2) the rate of information perishing. The optimal recovery period decreases (buying more information), if either of these factors increases.

The effect of risk aversion on the optimal information recovery policies was studied briefly. The optimality condition for an expected-value decision maker was extended to the case of a risk-averse decision maker. It was found very difficult to make any comparison between the two cases, however, because the payoff lotteries are different for the two decision makers. For a simple Markovian state and a quadratic payoff function, the optimal information recovery periods for an expected-value and a risk-averse decision maker were calculated and compared for various values of the parameters involved. In all cases the information recovery period for a risk-averse decision maker was found larger than the one for the expected-value decision maker. This suggests that a risk-averse decision maker buys less information than an expected-value decision maker, when there is uncertainty regarding the time of the decision. Whether (or under what conditions) this is in fact true, remains uncertain.

CHAPTER 5

A POSTERIORI OPTIMAL INFORMATION RECOVERY

In this chapter the a posteriori policies for recovery of information are investigated. For this type of policy, as mentioned in Chapter 3, the result of the previous observations is used in deciding the time of the next observation(s). The optimality condition is found to be similar to that of the a priori policies. We will define the "value of new information" which is found to be more appropriate than "the residual value of past information" for studying the a posteriori policies. Using the new definition we will find conditions under which the two types of policies coincide.

5.1 Optimality Conditions for A Posteriori Policies

It is easily seen that the same argument which was used in Chapter 4 (Sec. 4.2) to obtain the optimality condition for the a priori case can be used here, except that all the a priori payoffs must be replaced by the a posteriori payoffs which depend on the result of the previous observations. For a Markovian system and a perfect observation the only information needed at each point of time is the result of the last observation of the system. Consquently, the computations are greatly reduced in this case. Our results in this chapter are general, however, and do not require the Markovian property. If Z_o is the result of all the previous observations (up to the last observation at time zero), the optimality condition (4.15), when written for the a posteriori case is,

$$V_t(t_1^*, Z_o) = V_t(t_1^* + 1, Z_o) \tag{5.1}$$

AD-A074 956 STANFORD UNIV CALIF DEPT OF ENGINEERING-ECONOMIC SYSTEMS F/G 12/1

OPTIMAL INFORMATION ACQUISITION FOR RANDOMLY OCCURRING DECISION--ETC(U)

SEP 79 A SHARIFNIA N00014-79-C-0036

UNCLASSIFIED EES-DA-79-1 NL

2 OF 2

AD A 074956

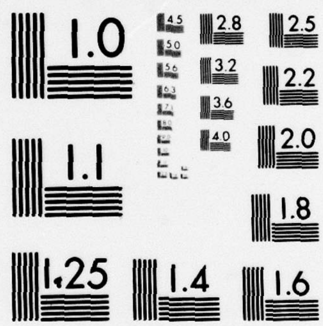

MICROCOPY RESOLUTION TEST CHART

where

$$V_t(t_1, Z_o) = \text{maximum net expected payoff from time } t \text{ on, with the next observation at } t_1 (t_1 \geq t) \text{, given } Z_o . \tag{5.2}$$

Note that $\bar{V}_t(t_1)$ defined for the a priori case is the expected value of $V_t(t_1, Z_o)$ over all values of Z_o. From (5.1) the necessary condition for optimum recovery time (t^*) can be obtained in exactly the same way as in the a priori case. We get

$$V(t^*, Z_o) = g_{t^*} \cdot R(t^*, Z_o) + (1-g_{t^*}) \cdot V(t^*+1, Z_o) \tag{5.3}$$

where $V(t, Z_o) = V_t(t, Z_o)$ (defined in (5.2)), and

$R(t, Z_o)$ = residual value of information at time t given that the result of the previous observations was Z_o.

Equation (5.3) has exactly the same form as the optimality condition for the a priori case (Eq. (4.17)), but with a posteriori payoffs substituted for the a priori ones. Consequently the same interpretation as in the a priori case holds for this case, namely that t^* is such that the marginal expected loss at present (t^*) of buying information slightly later $(t^*+1$ rather than $t^*)$ is equal to its marginal benefit in the future (from t^*+1 on). The benefits and losses, however, are a posteriori ones, depending on the result of the previous observations, Z_o.

5.2 Calculation of $t^*(Z_o)$

If t^* depends on Z_o , there is no easy way to find $t^*(Z_o)$ in general. The optimality condition found may be helpful but is not often sufficient for calculating $t^*(Z_o)$. For the finite-horizon case we may use dynamic programming to find $t^*(Z_o)$. For the infinite-horizon case the method of policy iteration [7] may be used. Both cases are briefly explained below.

1. Finite Horizon: To use the backward dynamic programming method we have to change our definitions slightly. Let us define:

$V_t(Z_{t-\tau})$ = maximum expected payoff from time t on, given $Z_{t-\tau}$;

$R_t(Z_{t-\tau})$ = expected residue of information at time t , given $Z_{t-\tau}$;

$V'_t(Z_{t-\tau})$ = expected payoff at time t with an observation at t , given $Z_{t-\tau}$.

The decision in each period is whether or not to buy new information at that time,as shown in Fig. 5.1.

t, buy inf.: $$g_t \cdot V'_t(Z_{t-\tau}) + (1-g_t) \cdot \langle V_{t+1}(Z_t) | Z_{t-\tau}, \mathscr{B} \rangle - c = W_1(t)$$

t, do not buy inf.: $$g_t \cdot R_t(Z_{t-\tau}) + (1-g_t) \cdot V_{t+1}(Z_{t-\tau}) = W_2(t)$$

Figure 5.1. Desision at time t in dynamic programming formulation.

$W_1(t)$ and $W_2(t)$ are the expected payoffs (at t) with and without an observation at t , respectively. We have

$$V_t(Z_{t-\tau}) = \max\{W_1(t) , W_2(t)\} \tag{5.4}$$

At horizon N we have

$$V_N(z_\tau) = \max\{g_N \cdot R_N(Z_{N-\tau}) , g_N \cdot V'_N(Z_{N-\tau})-c\}$$

Working backward we can calculate $V_t(Z_{t-\tau})$ for $t = N-1, N-2,\ldots,1,0$. The burden of computation increases rapidly as N or the number of possible values of Z increases.

2. Infinite Horizon: For the infinite-horizon case the method of policy iteration may be used to obtain $t^*(Z_o)$. We can use our optimality condition (Eq. (5.3)) to simplify the calculations. The optimality condition is:

$$V(t^*,Z_o) = g_{t^*} \cdot R(t^*,Z_o) + (1-g_{t^*}) \cdot V(t^*+1,Z_o) \tag{5.5}$$

Also we have

$$V(0,Z_o) = \sum_{t=o}^{T^*-1} p_t \cdot R(t,Z_o) + P_{t \geq t^*} \cdot V(t^*,Z_o) - c \tag{5.6}$$

This condition follows from the optimality of t^* and splits the maximum expected payoff from time zero on $(V(0,Z_o))$ into the expected payoff from 0 to t^*-1 and the expected payoff from t^* on $(V(t^*,Z_o))$. p_t is the probability that the decision will occur at time t and $P_{t \geq t^*}$ is the probability that it will occur at or after t^* . $t^*(Z_0)$ can be found by the following procedure:

(1) Select (guess) a policy $t^* = f^1(Z_o)$, calculate $V^1(0,Z_o)$ from (5.6) by setting $V(t^*,Z_o)$ at an arbitrary value, and iterate until convergence.

(2) Calculate $V^1(t^*,Z_o)$ and $V^1(t^*+1,Z_o)$ from $V^1(0,Z_o)$ and substitute in Eq. (5.5) to find the improved policy $t^* = f^2(Z_o)$.

(3) Continue (1) and (2) until $t^* = f^i(Z_o)$ converges to the optimal policy $t^* = \hat{f}(Z_o)$.

This method, although useful, does not guarantee the convergence to the optimum for our problem and often convergency problems arise. Fortunately, for an important class of problems t^* is not a function of Z_o and is relatively easy to calculate. This class of problems, however, is not easily distinguished in our present formulation, in which the "residual value of past information" is the main variable. In the following we will introduce a new variable which may be thought of as a dual to the residual value of past information. This variable is found to be much more appropriate for the study of the a posteriori recovery of information. In particular, using this variable makes it easier to distinguish the class of problems for which t^* does not depend on Z_o .

5.3 Value of New Information

The new variable is the "value of new information" at each point of time (given our state of information at that time), and is shown in Fig. 5.2. $V_1(t,Z_o) = R(t,Z_o)$ is the expected payoff of the decision at

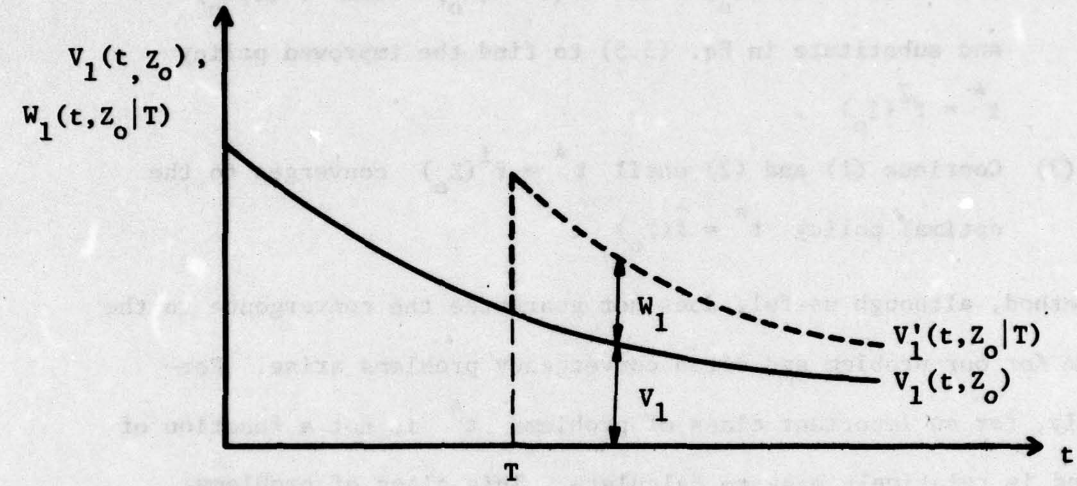

Figure 5.2. Value of new information $W_1(t,Z_o|T)$.

time t, given that Z_o was observed at the previous observations of the state. $V_1'(t,Z_o|T)$ denotes the expected payoff of the decision at time t with new information at T, given Z_o. The difference

$$W_1(t, Z_o|T) = V_1'(t,Z_o|T) - V_1(t,Z_o) \tag{5.7}$$

is the expected value of the new information (at T) for making the decision at time t. We will substitute $W_1(t,Z_o|T)$ for $V_1(t,Z_o)$ in our formulations. Therefore, rather than seeing how valuable our information is from the past $(V_1(t,Z_o))$, we see how valuable the new

information is at each point of time $(W_1(t,Z_o|T))$. In addition to simplifying the optimality condition, the new variable has the advantage that while $V_1(t,Z_o)$ and $V_1'(t,Z_o|T)$ always depend on Z_o , the difference $W_1(t,Z_o|T)$ does not depend on Z_o for an important class of problems. As we will see in Section 5.5, this has the important implication that t^* is independent of Z_o for this class of problems. Working with $W_1(t,Z_o|T)$ we can immediately see whether or not t^* depends on Z_o . This is not the case if we work with $V_1(t,Z_o)$ and we may get involved in difficult computations even if t^* is not a function of Z_o, and therefore can be found in much easier ways.

5.4 Optimality Condition in Terms of "Value of New Information"

The necessary condition for optimality (Eq. (5.3)) can be written in terms of the value of new information. Let us define $W(t,Z_o)$ as

$$W(t,Z_o) = V(t,Z_o) - \sum_{t'=t}^{N} p_{t'} \cdot V_1(t',Z_o) \tag{5.8}$$

Recall that $V(t,Z_o)$ was the maximum expected payoff from time t on (with recovery at t) , given Z_o . The summation on the right-hand side is the net expected payoff from t on, if no more information is bought ($p_{t'}$ is the probability that the decision will occur at t' , and N is the horizon). Therefore $W(t,Z_o)$ is the net value of all future purchases of information (including the first one at time t) . Notice that $W(t,Z_o)$ is the counterpart of $V(t,Z_o)$ in the new formulation. Solving Eq. (5.8) for $V(t,Z_o)$ and substituting in the optimality condition (5.3), we have

$$W(t^*,Z_o) + \sum_{t'=t^*}^{N} p_{t'} \cdot R(t,Z_o) = g_{t^*} R(t^*,Z_o) + (1-g_{t^*}) \cdot \left[W(t^*+1,Z_o) + \sum_{t'=t^*+1}^{N} p_{t'} \cdot R(t',Z_o) \right]$$

By subtracting $\sum_{t'=t^*}^{N} p_{t'} \cdot R(t',Z_o)$ from both sides we find

$$W(t^*,Z_o) = (1-g_{t^*}) \cdot W(t^*+1,Z_o) \tag{5.9}$$

This equation is the necessary condition for optimality written in terms of the value of new information and has a much simpler form than Eq. (5.3).

5.5 Independence of the Optimal Revovery Time (t^*) from the Result of the Previous Observations (Z_o)

As mentioned earlier, if t^* depends on the realization of the previous observations (Z_o), it is often very difficult to calculate. We also mentioned that, fortunately, for an important class of problems t^* does not depend on Z_o and is therefore relatively easy to calculate. In this section we will find conditions under which t^* is independent of Z_o. We will then identify the mentioned class of problems.

__Theorem 5.1__. If the (total) expected value of new information when bought only once is always independent of the realization of the previous observations (Z_o), then the optimal recovery time is independent of Z_o.

__Proof__: The (total) expected value of new information, when bought only once, is the net expected gain from the time of the information recovery (T) up to the horizon (N) assuming that no more information

is bought in this interval (Fig. 5.3). Denoting this value by $W^1(T,Z_o)$

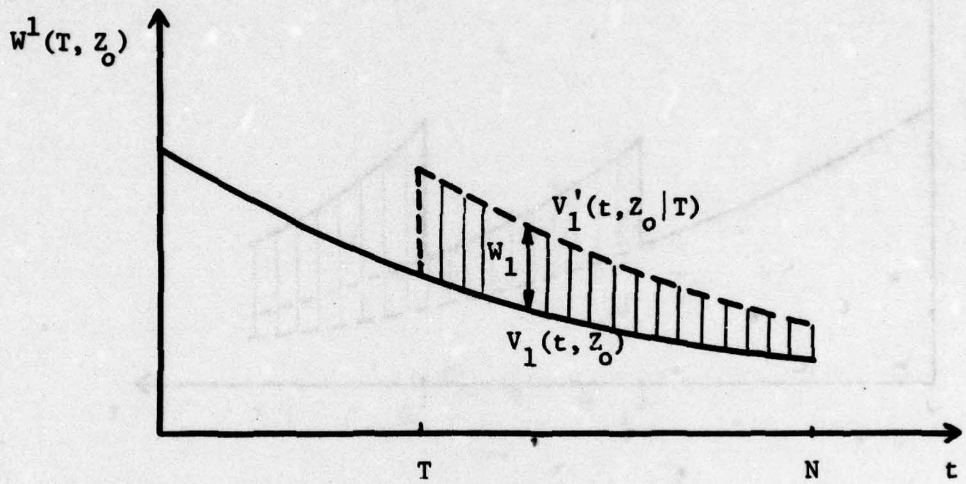

Figure 5.3. Value of new information when bought only once.

we have

$$W^1(T,Z_o) = \sum_{t=T}^{N} p_t \cdot W_1(t,Z_o|T) - c$$

where p_t is the probability of the occurrence of the decision at time t and c is the cost of the information. If information is recovered more than once (Fig. 5.4), the expected benefit ($W(0)$) is the sum of the

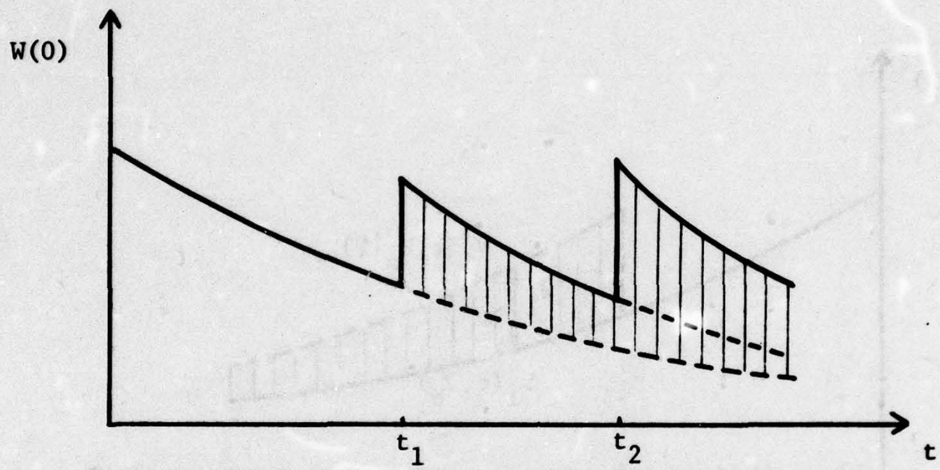

Figure 5.4. Value of future purchases of information.

expected value of each piece of information when bought alone over all information recoveries:

$$W(0) = \sum_i p_{t \geq t_i} \cdot W^1(t_i, Z_{i-1}) \tag{5.10}$$

where $p_{t \geq t_i}$ is the probability that the decision occurs after the i^{th} recovery. The information available at time t_i is $Z_{t_{i-1}}$, which is shown, for simplicity, by Z_{i-1}. Now suppose that $\hat{t}_i(Z_{i-1})$ is the optimal time for the i^{th} recovery as a function of the realization of the previous observations. The maximum expected value (at $t=0$) of all future information recoveries is

$$W(0,Z_o) = \left\langle \sum_i P_{t \geq t_i} \cdot W^1(\hat{t}_i(Z_{i-1}), Z_{i-1} | Z_o, \mathscr{E} \right\rangle \tag{5.11}$$

If the value of new information when bought only once $(W^1(t,Z))$ does not depend on the realization of the previous observations (but on t only), Eq. (5.11) can be written

$$W(0,Z_o) = \left\langle \sum_i P_{t \geq t_i} \cdot W^1(\hat{t}_i(Z_{i-1})) \middle| Z_o, \mathscr{E} \right\rangle \tag{5.12}$$

Now suppose that the information recovery schedule $\{t_1^*, t_2^*, \ldots\}$ maximizes $W(t_1, t_2, \ldots)$ defined as

$$W(t_1, t_2, \ldots) = \sum_i P_{t \geq t_i} \cdot W^i(t_i)$$

then we have

$$\sum_i P_{t \geq t_i} \cdot W^1(\hat{t}_i(Z_{i-1})) \leq W(t_1^*, t_2^*, \ldots) \tag{5.13}$$

and since (5.13) is true regardless of values of Z_i, from (5.12) it follows that

$$W(0,Z_o) \leq W(t_1^*, t_2^*, \ldots) \tag{5.14}$$

but $W(t_1^*, t_2^*, \ldots)$ is the expected benefit without using the result of observations (Z_i) in deciding t_i (or namely the maximum expected payoff of the a priori policy). It cannot be greater than the

expected payoff of the a posteriori policy, namely $W(0,Z_o)$. Therefore, from (5.14) we must have

$$W(0,Z_o) = W(t_1^*, t_2^*, \ldots) .$$

It follows, therefore, that the schedule $\{t_1^*, t_2^*, \ldots\}$ is also optimum for the a posteriori information recovery. Therefore, $\hat{t}_i(Z_{i-1}) = t_i^*$ is independent of Z_{i-1}.

Theorem 5.2. For a quadratic payoff function, if the state $\underline{s}(t)$ and the observation $\underline{z}(t)$ have normal distribution, then the optimal information recovery period is independent of the realization of the previous observations.

Proof: We show that for a quadratic payoff function and a normal state (and observation), the expected value of new information, when bought only once, is independent of the realization of the previous observations (Z_o). It then follows from Theorem 5.1 that the optimal recovery time is independent of Z_o. For a quadratic payoff function the expected payoff was found in Chapter 2. We showed that (Eq. (2.10))

$$V_1(\underline{y}) = \underline{\tilde{s}}'(\underline{y}) \cdot M \cdot \underline{\tilde{s}}(\underline{y}) \tag{5.15}$$

where $M = -1/2\, GH^{-1}G'$ (G and H are coefficient matrices of the payoff function, Eq. (2.5)), and $\underline{\tilde{s}}(\underline{y})$ is the posterior mean of the state vector $\underline{s}$, given information $\underline{y}$:

$$\underline{\tilde{s}}(\underline{y}) = \langle \underline{s} \mid \underline{y}, \mathscr{E} \rangle \tag{5.16}$$

In Fig. 5.5 $V_1(t,Z_o)$ is the expected payoff at time t given Z_o. $V_1'(t,Z_o|T)$ is the expected payoff at time t $(t > T)$ with new information

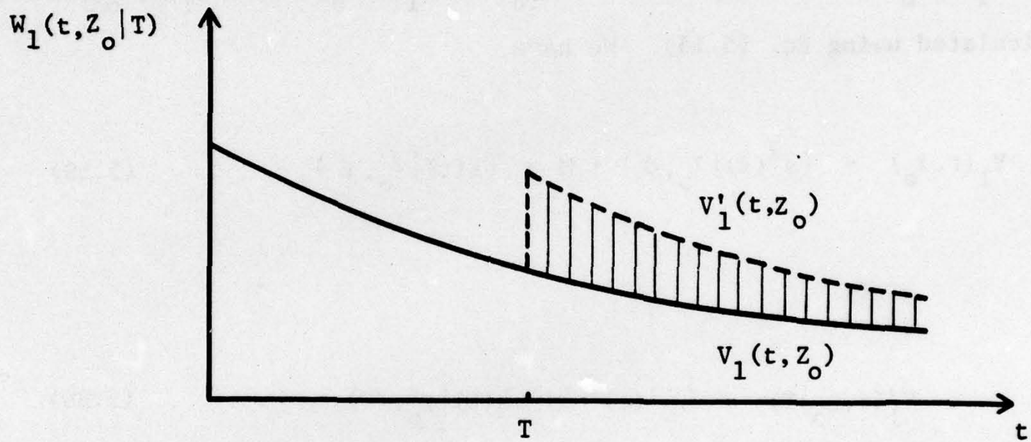

Figure 5.5 Value of new information when bought only at T .

at T , given $\underline{Z}_o$. The total expected gain of buying new information is:

$$W(T,\underline{Z}_o) = \sum_{t \geq T} p_t \cdot W_1(t,Z_o|T) - c \tag{5.17}$$

where

$$W_1(t,Z_o|T) = V_1'(t,Z_o|T) - V_1(t,Z_o) \tag{5.18}$$

p_t is the probability that the decision will occur at time t and c is the cost of the information. From Eq. (5.18) we can see that if $W_1(t,Z_o|T)$ is always (for all t and T) independent of Z_o, then $W(T,Z_o)$ will be independent of Z_o. Therefore it is sufficient to show that $W_1(t,Z_o|T)$ is independent of Z_o. $V_1(t,Z_o)$ and $V_1'(t,Z_o|T)$ can be calculated using Eq. (5.15). We have

$$V_1(t,Z_o) = \langle \underline{s}'(t)|Z_o,\mathscr{E}\rangle \cdot M \cdot \langle \underline{s}(t)|Z_o,\mathscr{E}\rangle \tag{5.19}$$

and

$$V_1'(t,Z_o|T) = \langle \tilde{\underline{s}}'(t) \cdot M \cdot \tilde{\underline{s}}(t)|Z_o,\mathscr{E}\rangle \tag{5.20}$$

where

$$\tilde{\underline{s}}(t) = \langle \underline{s}(t)|\underline{z}_T,Z_o,\mathscr{E}\rangle$$

Equation (5.20) can be written as

$$V_1(t,Z_o|T) = \mathrm{tr}(M\cdot\Sigma_{\tilde{\underline{s}}}) + \langle \tilde{\underline{s}}'(t)|Z_o,\mathscr{E}\rangle\cdot M \cdot \langle \tilde{\underline{s}}(t)|Z_o,\mathscr{E}\rangle \tag{5.21}$$

where $\Sigma_{\tilde{\underline{s}}}$ is the matrix of the covariances of $\tilde{\underline{s}}(t)$:

$$\Sigma_{\tilde{\underline{s}}} = \mathrm{cov}\langle \tilde{\underline{s}}(t)|Z_o,\mathscr{E}\rangle$$

Noting that

$$\langle \tilde{\underline{s}}(t) | Z_o, \mathscr{E} \rangle = \langle \langle \underline{s}(t) | \underline{z}_T, Z_o, \mathscr{E} \rangle | \underline{Z}_o, \mathscr{E} \rangle$$

$$= \langle \underline{s}(t) | Z_o, \mathscr{E} \rangle \tag{5.22}$$

and in view of (5.19), (5.21) can be written as:

$$V_1'(t, Z_o | T) = tr(M \cdot \Sigma_{\tilde{\underline{s}}}) + V_1(t, Z_o) .$$

Therefore

$$W_1(t, Z_o | T) = V_1'(t, Z_o | T) - V_1(t, Z_o) = tr(M \cdot \Sigma_{\tilde{\underline{s}}}) \tag{5.23}$$

Now we show that if $\underline{s}(t)$ and $\underline{z}(t)$ have normal distribution, $\Sigma_{\tilde{\underline{s}}}$ is independent of Z_o . To show this, we can use the following relation which decomposes the prior variance into the variance of the posterior mean and the expected value of the posterior variance [16] :

$$^V\langle \underline{x} | \mathscr{E} \rangle = {}^V\langle \langle \underline{x} | \underline{y}, \mathscr{E} \rangle | \mathscr{E} \rangle + \langle {}^V\langle \underline{x} | \underline{y}, \mathscr{E} \rangle | \mathscr{E} \rangle .$$

Letting $\underline{x} = \underline{s}(t)$, $\underline{y} = \underline{z}_T$ and $\mathscr{E} = (Z_o, \mathscr{E})$, we have

$$^V\langle \underline{s}(t) | Z_o, \mathscr{E} \rangle = {}^V\langle \langle \underline{s}(t) | \underline{z}_T, Z_o, \mathscr{E} \rangle | Z_o, \mathscr{E} \rangle + \langle {}^V\langle \underline{s}(t) | \underline{z}_T, Z_o, \mathscr{E} \rangle | Z_o \mathscr{E} \rangle$$

but the first term on the right-hand side is $\Sigma_{\underline{\tilde{s}}}$; therefore,

$$\Sigma_{\underline{\tilde{s}}} = {}^{V}\langle \underline{s}(t) | Z_o, \mathscr{E} \rangle - \langle {}^{V}\langle \underline{s}(t) | \underline{z}_T, Z_o, \mathscr{E} \rangle | Z_o, \mathscr{E} \rangle$$

If $\underline{s}(t)$ and $\underline{z}(t)$ have normal distribution, then both ${}^{V}\langle \underline{s}(t) | Z_o, \mathscr{E} \rangle$ and ${}^{V}\langle s(t) | \underline{z}_T, Z_0, \mathscr{E} \rangle$ are independent of Z_o.(*) Therefore $\Sigma_{\underline{\tilde{s}}}$ is independent of Z_o and the proof is complete.

5.6 A Posteriori Recovery of Information for the Bidding Example

The a priori recovery of information for the Bidding Example of Chapter 2 was studied in Chapter 4. Here we will find the a posteriori policy for the recovery of information for that example. The uncertain variable of interest in the Bidding Example was the cost of performing the bidding contract (p) . We assumed that p has a constant mean over time, and that its variation from the mean ($\Delta P = s$) changes over time according to the linear Markovian system

$$s(t) = \lambda \cdot s(t-1) + \varepsilon(t) \tag{5.24}$$

We have also assumed that our observation of s is perfect. Consequently, the only information needed at each time is the result of the last observation of s (s_o) . To find the optimum a posteriori policy for updating our information about s(t) , let us first discover whether such a policy would depend on the result of the previous observation (s_o) . By Theorem 5.1 if the value of new infomation when bought only once is

(*) See updating relations for normal random variables in [2], for example.

independent of the realization of the previous observation (s_o), then the optimal time for the next recovery (t^*) is independent of s_o. The expected gain at time t of buying new information at T is

$$W_1(t,s_o|T) = V_1'(t,s_o|T) - V_1(t,s_o)$$

$W_1(t,s_o|T)$ is shown in Fig. 5.6.

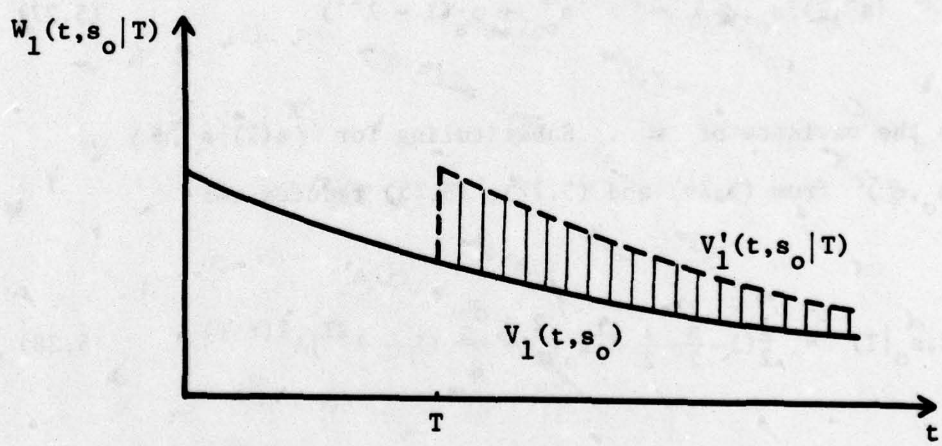

Figure 5.6. Value of new information bought at T, given s_o

$V_1'(t,s_o|T)$ is

$$V_1'(t,s_o|T) = \langle V_1(t-T,s(T))|s_o,\mathscr{E}\rangle .$$

Substituting for $V_1(t-T, s(T))$ from Eq. (2.48), we find

$$V_1'(t,s_o|T) = \frac{1}{2}\left\langle(1 - \frac{m}{2} - \frac{1}{2}\lambda^{t-T}s(T))^2 | s_o, \mathscr{E}\right\rangle$$

$$= \frac{1}{2}\left[(1 - \frac{m}{2})^2 - (1 - \frac{m}{2})\lambda^{t-T}\langle s(T)|s_o,\mathscr{E}\rangle + \frac{1}{4}\lambda^{2(t-T)}\langle s^2(T)|s_o,\mathscr{E}\rangle\right] \tag{5.25}$$

From (5.24) it is easy to show that

$$\langle s(T)|s_o,\mathscr{E}\rangle = \lambda^T s_o \tag{5.26}$$

and

$$\langle s^2(T)|s_o,\mathscr{E}\rangle = \lambda^{2T}s_o^2 + \sigma_s(1 - \lambda^{2T}) \tag{5.27}$$

where σ_s is the variance of s . Substituting for $\langle s(T)|s_o,\mathscr{E}\rangle$ and $\langle s^2(T)|s_o,\mathscr{E}\rangle$ from (5.26) and (5.27), (5.25) reduces to:

$$V_1'(t,s_o|T) = \frac{1}{2}(1-\frac{m}{2}-\frac{1}{2}\lambda^t s_o)^2 + \frac{\sigma_s}{8}(1 - \lambda^{2T})\lambda^{2(t-T)} \tag{5.28}$$

From this equation and Eq. (2.48) we have

$$W_1(t,s_o|T) = V_1'(t,s_o\ T) - V_1(t,s_o)$$

$$= \frac{\sigma_s}{8}(1 - \lambda^{2T})\lambda^{2(t-T)} \tag{5.29}$$

Notice that both $V_1(t,s_o)$ and $V_1'(t,s_o|T)$ depend on s_o but their difference (value of new information at t) does not depend on s_o .

It follows that the value of new information when bought only once is independent of s_o and, therefore, t^* is independent of s_o. Equivalently, the a priori and the a posteriori information recovery policies are the same. Therefore the periodic information recovery policy found for the a priori case is also optimum for this case. The optimum information-recovery period was found to be (Eq. (4.25)):

$$T^* \cong \frac{\log \frac{W}{\sigma_s/8}}{\log \frac{1}{\rho}}$$

where W is the (total) value of new information when bought only once (given that the past information is completely perished) and ρ is the rate of information perishing. Note that the value of new information when bought only once appears as an important variable in the study of the optimal information recovery policies.

5.7 Bounds on $t^*(Z_o)$

It was mentioned in Section 5.2 that if t^* depends on Z_o, it is often difficult to calculate. In this section we find upper and lower bounds for $t^*(Z_o)$. These bounds are easily obtained from the notion of "value of new information" and are specially useful when t^* depends on Z_o.

(1) Lower bound on $t^*(Z_o)$: Consider one purchase of information. The expected benefit from this information is maximum, if no more information

is bought in the future (because future information will demolish the residual value of this information). This maximum benefit must be greater than the cost of information (c). In other words, the net value of information, when bought only once, must be positive:

$$W^1(t^*,Z_o) = \sum_{t=t^*}^{N} p_t \cdot W_1(t,Z_o|t^*) - c \geq 0 \qquad (5.30)$$

Here $W_1(t,Z_o|t^*)$ is the gain at time t of buying new information at t^* (see Fig. 5.6) and p_t is the probability that the decision will occur at time t. If $W^1(t,Z_o)$ is increasing with t (which is often the case, especially if we are not too close to the horizon), then $t_1(Z_o)$, satisfying the equation

$$W^1(t_1,Z_o) = 0 \qquad (5.31)$$

is a lower bound on $t^*(Z_o)$.

(2) Upper bound on $t^*(Z_o)$: Considering again one purchase of information (say at t), the expected benefit from this information is minimum if it is intended for use at t only (this is the case when, for instance, the next purchase will be at $t+1$). This minimum value cannot be greater than the cost of information (because if the expected benefit at t exceeds the cost of information, we cannot be worse off by buying it at t rather than at a later time). Therefore we have

$$p_{t^*} \cdot W_1(t^*, Z_o | t^*) \leq c \quad . \tag{5.32}$$

Assuming that $p_t \cdot W_1(t, Z_o | t)$ is increasing with t (which is very often the case), then $t_2(Z_o)$ satisfying the equation

$$p_{t_2} \cdot W_1(t_2, Z_o | t_2) = c \tag{5.33}$$

is an upper bound on $t^*(Z_o)$.

Example 5.1 We have assumed so far that the payoff function is quadratic. We have shown that for a quadratic payoff function and a normal state t^* is independent of Z_o (Theorem 5.2). Here we give an example of a nonquadratic payoff function for which t^* depends on Z_o (for a linear Markovian system). Then we use the results of this section to find lower and upper bounds on $t^*(Z_o)$. Let us consider the following payoff function:

$$v_1(s,d) = s^2 d - \frac{1}{2} d^2 \tag{5.34}$$

Maximizing $v_1(s,d)$ is equivalent to minimizing $(s^2 - d)^2$. Assuming that $s(t)$ changes according to the linear Markovian system,

$$s(t) = \lambda\, s(t-1) + \varepsilon(t)$$

and has a normal distribution, we can show that

$$V_1(t,s_o) = \frac{1}{2}\left[\lambda^{2t} s_o^{\,2} + \sigma_s(1-\lambda^{2t})\right]^2$$

and

$$V_1'(t,s_o|T) = \frac{1}{2}\left[\lambda^{2t} s_o^{\,2} + \sigma_s(1-\lambda^{2t})\right]^2$$
$$+ \sigma_s(1-\lambda^{2T})\left[2\lambda^{2T} s_o^{\,2} + \sigma_s(1-\lambda^{2T})\right]\cdot \lambda^{4(t-T)}$$

Therefore the expected benefit at time t from buying new information at T is:

$$W_1(t,s_o|T) = V_1'(t,s_o|T) - V_1(t,s_o)$$
$$= \sigma_s(1-\lambda^{2T})\left[2\lambda^{2T} s_o^{\,2} + \sigma_s(1-\lambda^{2T})\right]\cdot \lambda^{4(t-T)} \tag{5.35}$$

Note that $W_1(t,s_o|T)$ depends on s_o, and therefore we expect that t^* will be a function of s_o too. In the following we will find upper and lower bounds for $t^*(s_o)$.

(1) Upper bound on $t^*(s_o)$: From Eq. (5.33) we have (assuming $g_t = g$ to be constant) :

$$g \cdot W_1(t_2,s_o|t_2) = c \tag{5.36}$$

Substituting for $W_1(t_2,s_o|t_2)$ from (5.35) we have

$$g \cdot \sigma_s(1-\lambda^{2t_2})\left[2\lambda^{2t_2} s_o^2 + \sigma_s(1-\lambda^{2t_2})\right] = c \tag{5.37}$$

Solving (5.37) for λ^{2t_2}, we find

$$\lambda^{2t_2} = \frac{\sigma_s - s_o^2}{\sigma_s - 2s_o^2} \pm \sqrt{\left(\frac{\sigma_s - s_o^2}{\sigma_s - 2s_o^2}\right)^2 - \frac{\sigma_s - c/g\sigma_s}{\sigma_s - 2s_o^2}} \tag{5.38}$$

where t_2 is an upper bound for $t^*(s_o)$.

(2) Lower bound on $t^*(s_o)$: From Eq. (5.31) we have

$$\sum_{t=t_1}^{\infty} p_t \cdot W_1(t,s_o|t_1) = c$$

Substituting for $W_1(t,s_o|t_1)$ from Eq. (5.35) we find:

$$W_1(t_1,s_o|t_1) \cdot \sum_{t=t_1}^{\infty} g\left((1-g)\lambda^4\right)^{t-t_1} = c$$

or

$$\frac{g}{1-(1-g)\lambda^4} \cdot W_1(t_1,s_o|t_1) = c \tag{5.39}$$

Comparing this equation with that of the upper bound (5.36) we notice that they are identical, except that g in (5.36) is replaced by $u = g/(1-(1-g)\lambda^4)$ in (5.39). Therefore t_1 can be obtained from (5.38) by replacing u for g in this equation. If we draw

$W_1(t,s_o|t)$ from Eq. (5.35), then from (5.36) and (5.39), t_2 and t_1 are the intersections of the horizontal lines of heights c/g and c/u with $W_1(t,s_o|t)$, respectively. This is shown in Fig. 5.7.

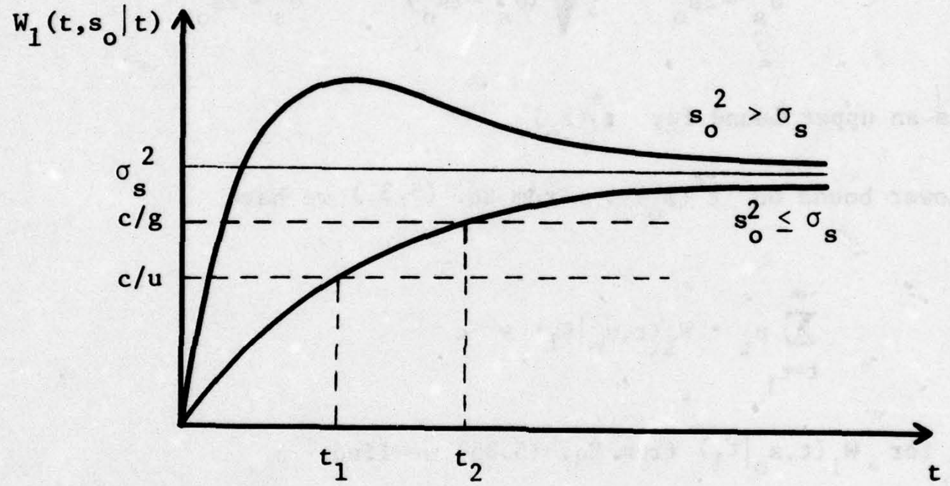

Figure 5.7. Upper and lower bounds for t^*, for a given value of s_o.

For $s_o^2 \le \sigma_s$, $W_1(t,s_o|t)$ always increases with t. Recall that this was necessary for the bounds to be valid. For $s_o^2 > \sigma_s$, $W_1(t,s_o|t)$ has an overshoot and does not always increase. Nevertheless, we can find the correct bounds by taking those intersections of the horizontal lines c/g and c/u with $W_1(t,s_o|t)$ which are in the increasing segment of $W_1(t,s_o|t)$. In Fig. 5.8 the upper and lower bounds as functions of s_o are depicted for numerical values of $g = .2$, $\lambda = .9$, $\sigma_s = 2.5$, and $c = 1$.

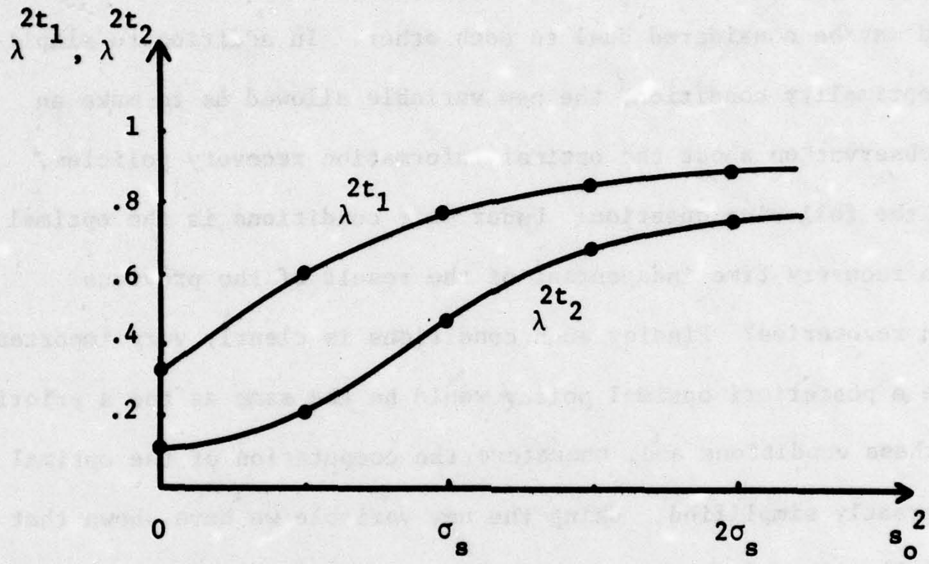

Figure 5.12. Upper and lower bounds as functions of s_o .

5.8 Summary

In this chapter the a posteriori optimal information recovery policies for a one-time contingent decision were investigated. The optimality condition for this case was shown to be the same as the optimality condition for the a priori case except that the a posteriori payoffs were substituted for the a priori ones. The optimal information recovery time depends, in general, on the result of the previous observations of the state.

In Chapter 4, where we studied the a priori information recovery policies, the main variable in the formulation of the optimality condition

was the "residual value of the past information" at each point of time. In this chapter we have defined a new variable, namely the "value of new information" at each point of time. The two variables are intimately related and may be considered dual to each other. In addition to simplifying the optimality condition, the new variable allowed us to make an important observation about the optimal information recovery policies, concerning the following question: Under what conditions is the optimal information recovery time independent of the result of the previous information recoveries? Finding such conditions is clearly very important, because the a posteriori optimal policy would be the same as the a priori one under these conditions and, therefore, the computation of the optimal policy is greatly simplified. Using the new variable we have shown that if the (total) value of the new information, when it is bought only once, is always independent of the result of the previous observations, then the optimal information recovery time is independent of the result of the previous observations. This result is general and assumes no conditions on the state, the decision, or the information structure. It is also relatively easy to check. An important example of the above property is when the payoff function is quadratic and the state and the observation have normal distributions.

Finally, when the optimal information recovery time depends on the result of the previous observations, we have found both upper and lower bounds for it. These bounds depend on the result of the previous observations.

CHAPTER 6

OPTIMAL INFORMATION RECOVERY FOR REPETITIVE CONTINGENT DECISIONS

In Chapters 4 and 5 we investigated the optimal informatior recovery policies for decisions which occur only once. In this chapter we will extend our study to the case where the decision may be repeated in time. For simplicity we will restrict ourselves to the a priori policies. We know from Chapter 5 that for an important class of problems the a priori and the a posteriori policies are the same.

6.1 Differences with the One-Time Decision Case

The most important difference with the one-time decision case is, of course, in the occurrence model. We assume that each occurrence of the decision is independent of the previous occurrences,

$$\{\text{Decision occur at } t|\mathscr{E}\} = g_t \tag{6.1}$$

For simplicity we assume throughout this chapter that the horizon is infinite and g_t is constant in time. There are two important differences between the two cases, concerning the recovery of information. In the repetitive case, (1) each piece of information may be used for more than one decision, and (2) there may be an opportunity to learn about the state of the system at the time of a decision from the outcome of that decision and use this information for later decisions. The first property results in buying more information compared to the one-time decision case because information is potentially more valuable in this case.

The second property, however, results in buying less information, since we obtain some free information from each decision. For the repetitive case we have to answer the following questions:

(1) What is T^*, namely the optimum recovery period assuming there is no interruption by decisions?

(2) How should T^* be revised when a decision occurs (and we receive some information from it)?

Clearly, the answers to these questions depend on the type of information learned from each decision. In the following we will answer the questions for four types of information learned from decisions: No information, perfect information, perfect but delayed information, and prompt but imperfect information.

Finally, since we assume infinite horizon we have to either discount the future payoffs (using the present-value criterion) or use other criteria (the rate of payoff, for example) to avoid the unboundedness of the total payoff.

6.2 Optimal Recovery of Information

I. No information from decisions.

This is an extreme case, where we learn nothing about the state from a decision. This might be the case when, for example, the outcome of the decision will not be known until after a very long delay. The necessary condition for the optimum recovery period T^* (the recovery period with no interruption by decisions) can be obtained by the same argument as in the one-time decision case (Section 4.2). From the result of that argument, T^* can be obtained by equating the expected payoff

with recovery at T^*, and at T^*+1, both calculated at time T^* (Eq. (4.15)). These expected payoffs are shown in Fig. 6.1. If information is

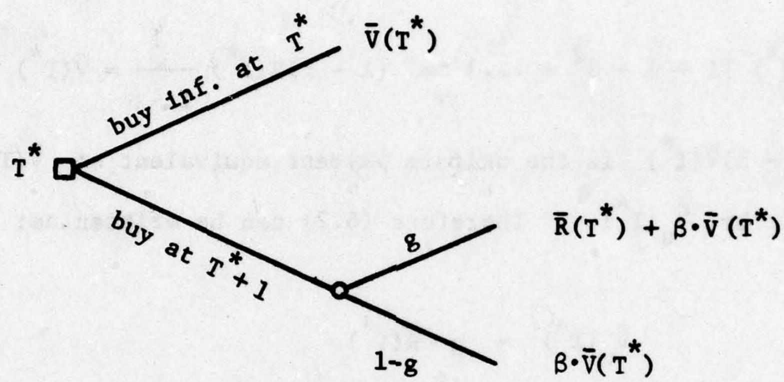

Figure 6.1. Payoffs at T^* with no information from decisions.

recovered at T^*, the maximum present value (at T^*) of all future payoffs is denoted by $\bar{V}(T^*)$. If information is recovered at T^*+1, then the maximum present value (at T^*) of all future payoffs is $\bar{R}(T^*) + \beta \cdot \bar{V}(T^*)$, if the decision happens at T^* ($\bar{R}(T^*)$ is the residual value of past information at T^*, and β is the discount rate), and $\beta\cdot\bar{V}(T^*)$, if the decision does not happen at T^*. Equating the expected payoffs of recovery at T^* and T^*+1, we have

$$\bar{V}(T^*) = g\cdot[\bar{R}(T^*) + \beta\cdot\bar{V}(T^*)] + (1-g)\cdot[\beta\cdot\bar{V}(T^*)] \quad .$$

After simplifying we find

$$(1 - \beta) \cdot \bar{V}(T^*) = g \cdot \bar{R}(T^*) \tag{6.2}$$

To find the interpretation of this condition, note that the present value of a sequence of uniform payments (paid each period) of amount $(1 - \beta)\bar{V}(T^*)$ is

$$(1 - \beta)V(T^*)\,[1 + \beta + \beta^2 + \ldots] = (1 - \beta)V(T^*)\frac{1}{1 - \beta} = \bar{V}(T^*)$$

Therefore, $(1 - \beta)\bar{V}(T^*)$ is the uniform payment equivalent of $\bar{V}(T^*)$, and we denote it by $\bar{V}_u(T^*)$. Therefore (6.2) can be written as:

$$\bar{V}_u(T^*) = g \cdot \bar{R}(T^*) \tag{6.3}$$

This condition states that the uniform payment equivalent of the optimal policy is equal to the expected payoff immediately before recovery. Comparing (6.3) with the optimality condition for the one-time decision case,

$$\bar{V}(T^*) = \bar{R}(T^*)$$

we can see a strong similarity between the two conditions. In the one-time decision case we can think that the decision will ultimately happen, and the expected payoff is equal to $\bar{R}(T^*)$. In the repetitive case, the uniform payment equivalent at each period is $\bar{R}(T)$ multiplied by the probability that the decision will occur at each period (g).

Since no information is learned from the decisions, the occurrence of a decision has no effect on the next recovery times and we will have a periodic recovery. Since in the repetitive case the result of each observation

may be used more than once to make the decision, the information is more valuable than in the one-time decision case. Consequently, we buy more information in the repetitive case and therefore, T^* for the repetitive case (with no information from decisions) is smaller than T^* for the one-time decision case.

II. Perfect information from decisions.

This case is the other extreme, namely when we learn perfect information about $s(t)$ from each decision. Since in this case the occurrence of a decision provides us with perfect information about the state, the past information will have no value after a decision occurs. As a result, each recovery of information has exactly the same value as in the case of the one-time decision (because in the one-time decision case the information is also used for making the decision once). It follows that T^* is the same as T^* in the one-time decision case. After each occurrence of the decision, however, we must revise the planned time for the next recovery. Since the occurrence of each decision is tantamount to a new observation, the next recovery time must be revised to T^* units after the time of the decision. Therefore the next recovery time is always T^* units from either the last recovery time or the last decision, whichever occurs later (Fig. 6.2).

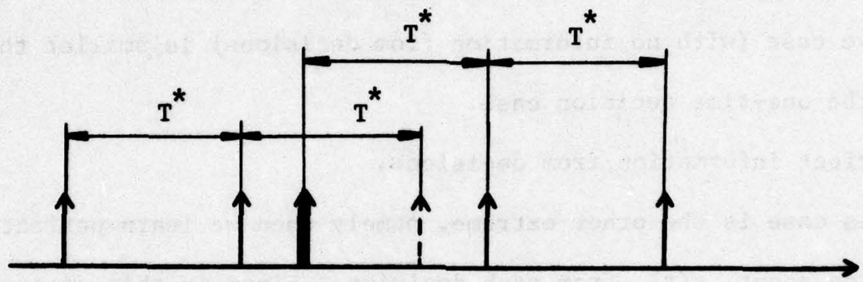

Figure 6.2. Optimal information recovery with perfect information from decision.

: Information recovery

: Occurrence of decision

Notice, however, that this policy maximizes the <u>rate of expected payoff</u> over time, but not necessarily the present value of the expected payoff. Let us find the optimality condition when the present-value criterion is used. By the same argument as before, we can find T^* by equating the expected payoffs (at T^*) of recovery at T^* and recovery at T^*+1. These payoffs are shown in Fig. 6.3. The present value (at T^*) of

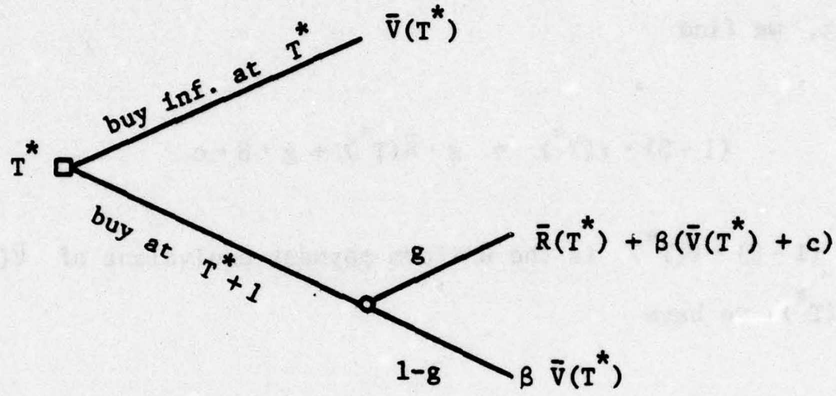

Figure 6.3. Expected payoffs at T^* with perfect information from decisions.

the expected payoff with recovery at T^* is denoted by $\bar{V}(T^*)$. Now consider the expected payoffs when we plan to buy information at T^*+1 . If the decision occurs at T^*, the expected payoff at T^* is $\bar{R}(T^*)$, and since we learn perfect information from this decision we will not buy information at T^*+1, and the expected future payoff is $\bar{V}(T^*) + c$ (c is the cost of information) because we obtained free information from the decision at T^*.† Therefore the total present value of the expected

---

†This is an approximation because $V(T^*) + c$ is the expected future payoff if information was learned at T^*+1 , while information is in fact learned at T^* .

payoff, if decision occurs at T^*, is $\bar{R}(T^*) + \beta(\bar{V}(T^*) + c)$. If the decision will not occur at T^*, then we buy information at $T^* + 1$ and the present value (at T^*) of the expected payoff is $\beta\bar{V}(T^*)$. Setting the present expected value of recovery at T^* and T^*+1 equal and simplifying, we find

$$(1-\beta) \cdot V(T^*) = g \cdot \bar{R}(T^*) + g \cdot \beta \cdot c$$

and since $(1-\beta) \cdot V(T^*)$ is the uniform payment equivalent of $\bar{V}(T^*)$, namely $\bar{V}_u(T^*)$, we have

$$\bar{V}_u(T^*) = g \cdot \bar{R}(T^*) + \beta \cdot g \cdot c \qquad (6.4)$$

Comparing this equation with the optimality condition for the case of no information from the decision (Eq. (6.3)), we notice that the benefits of learning $s(t)$ from each decision are represented by the term $\beta \cdot g \cdot c$. This term is the uniform payment equivalent of the benefits of learning from decisions. Intuitively, these benefits must be proportional to the cost of information (c), and the probability of the occurrence of the decision (g) . The reason $g \cdot c$ is discounted by β is that information from a decision (say at t) can be used for decisions from $t+1$ on (but not at t), while if the information is bought at t, it can be used from t on. Therefore information from the decision has an inherent one-unit delay and that is why $g \cdot c$ is discounted by β .

III. Perfect but delayed information from decisions.

In this case we assume that perfect information can be learned from each decision, but after a delay τ . Therefore, if the decision happens at time t , we will learn $s(t)$ at time $t+\tau$. This is the case when, for instance, the outcome of the decision will be revealed after a delay. Again we have to find T^* (the optimal recovery period with no interruptions by decisions), and the manner in which the planned recovery time must be revised after the occurrence of a decision. Let us define

t_i = time of the i^{th} occurrence of the decision after the last recovery,

T_i = optimal time for the next recovery after the i^{th} occurrence of the decision.

<u>Proposition 6.1</u>. If the time of receiving information from decision $(t_i + \tau)$ is before the planned time of the next recovery (T_{i-1}) , then the next recovery time must be revised to $t_i + T^*$. If the information from the decision is revealed after the planned time for the next recovery, then the next recovery time either remains unchanged or changes to $t_i + T^*$:

$$(1) \quad t_i + \tau \leq T_{i-1} \Longrightarrow T_i = t_i + T^*$$

$$(2) \quad t_i + \tau > T_{i-1} \Longrightarrow T_i = \begin{cases} T_{i-1} & \text{or} \\ t_i + T^* \end{cases}$$

Figure 6.4 illustrates the above statements .

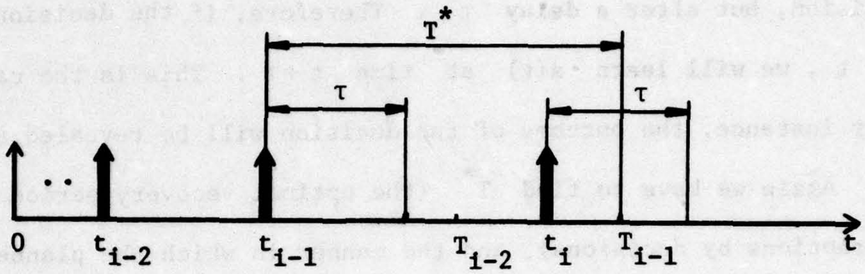

Figure 6.4. Revising the next recovery time after a decision:

$$T_{i-1} = t_{i-1} + T^* \; ; \quad T_i = \begin{cases} T_{i-1} & \text{or} \\ t_i + T^* \end{cases}$$

Proof: When $t_i + \tau \leq T_{i-1}$, the information from the decision is revealed to us before the planned time for the next recovery and we can benefit from this information. After the information is revealed, it is as if a recovery was made at t_i . Therefore the next recovery time must be revised to $t_i + T^*$.

(2) If $t_i + \tau > T_{i-1}$, the information from the decision will be revealed to us after the planned time for the next recovery. Therefore, if we want to obtain the benefits of this free information, we should not buy information at the planned time T_{i-1} . If we decide to do so, then t_i should be regarded as the last recovery time and therefore

$T_i = t_i + T^*$. There is an expected loss in so doing, however, because we will be low on information from T_{i-1} to $t_i + \tau$ (Fig. 6.5). If

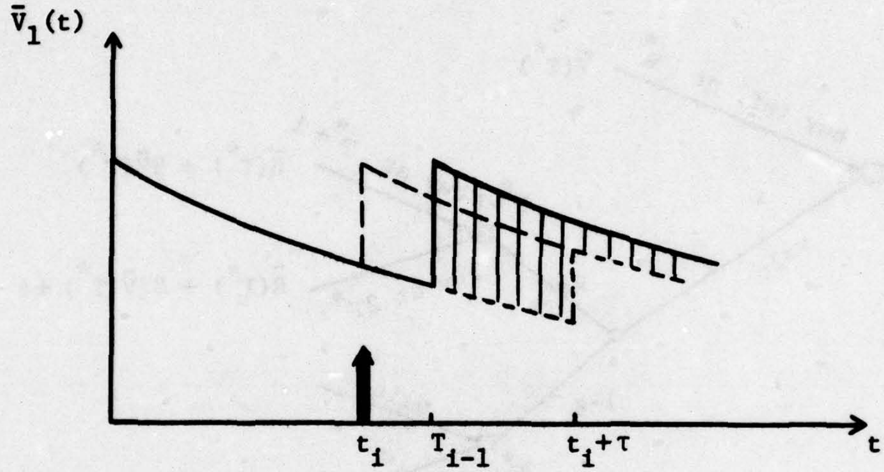

Figure 6.5. Loss due to not buying information at T_{i-1} .

this loss exceeds the benefits of the free information from the decision, we will be better off by ignoring the information from the decision and buying information at the planned time T_{i-1} . Therefore in this case T_i is either $t_i + T^*$ or T_{i-1} .

The optimality condition for T^* can be found, as before, by equating the present value (at T^*) of the expected payoffs with recovery at T^* and $T^* + 1$. However, since the recovery at $T^* + 1$ depends on whether or not the decision will occur at T^* , we must in

fact equate the present value of the expected payoffs (at T^*) of buying information and waiting as shown in Fig. 6.6. This figure is the same as Fig. 6.3 (perfect and prompt information from decision), except that

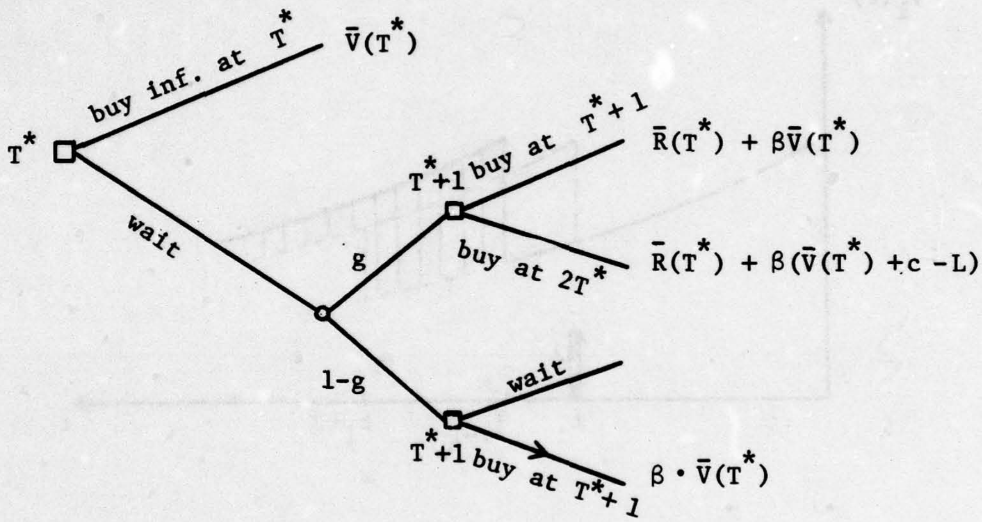

Figure 6.6. Expected payoffs with perfect but delayed information from the decisions.

if we do not buy information at T^*, and a decision occurs at T^*, then we have to decide whether to buy information as planned (at T^*+1) or to wait and benefit from the information from the decision at T^*. If we buy information at T^*+1 as planned, the present value of the expected payoff is $\bar{R}(T^*) + \beta V(T^*)$. If we wait, we must buy information at $T^* + T^* = 2T^*$. The present value of the expected payoff is

approximately[†] $\bar{R}(T^*) + \beta(\bar{V}(T^*) + c - L)$ where L is the present value of the loss (ℓ) by not buying information at $T^* + 1$, as shown in Fig. 6.7. Equating the present value of the expected payoff of buying

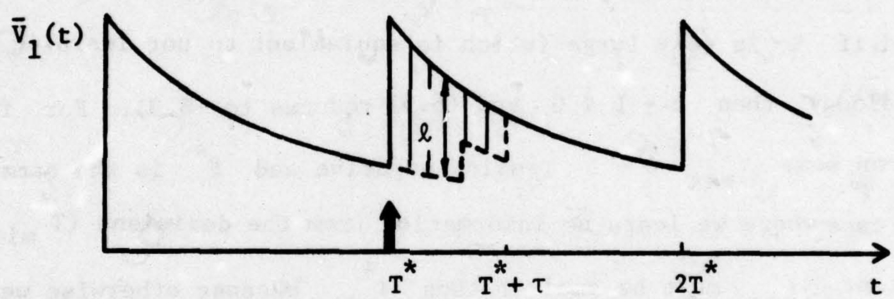

Figure 6.7. Loss of not buying information at T^*.

information at T^* with that of waiting, we have

$$V(T^*) = g \cdot \text{Max}\{\bar{R}(T^*) + \beta\bar{V}(T^*) \;,\; \bar{R}(T^*) + \beta(\bar{V}(T^*) + c - L)\}$$
$$+ (1-g)\beta\bar{V}(T^*)$$

After simplifying we find

$$\bar{V}_u(T^*) = g \cdot \bar{R}(T^*) + \beta \cdot g \,\text{Max}(0, c - L) \qquad (6.5)$$

[†]Ignoring the one-unit delay of information from the decision.

where $\bar{V}_u(T^*) = (1-\beta)\bar{V}(T^*)$ is, as before, the uniform payment equivalent of $\bar{V}(T^*)$. Comparing this equation with the optimality conditions of the previous cases (Eqs. (6.3) and (6.4)) we notice that the difference is in the second term on the right-hand side, which is the uniform payment equivalent of benefits of learning from the decisions. Moreover, (6.3) and (6.4) can be easily obtained from (6.5). If $\tau = 0$ (information prompt) then $L = 0$ and (6.5) is reduced to (6.4). On the other hand, if τ is very large (which is equivalent to not learning from decisions), then $c - L < 0$ and (6.5) reduces to (6.3). For τ greater than some τ_{max}, $c - L$ remains negative and T^* is the same as T^* in the case where we learn no information from the decisions (T^*_{min}). It is clear that τ_{max} must be smaller than T^*_{min} because otherwise we could not start using the information from the decisions. As τ decreases below τ_{max} T^* increases. For $\tau = 0$, T^* reaches its maximum (T^*_{max}), which is the same as T^* in the case of perfect and prompt information. T^* as a function of τ is shown in Fig. 6.8. Intuitively as τ decreases below τ_{max} we learn more from the decisions, and consequently we buy less

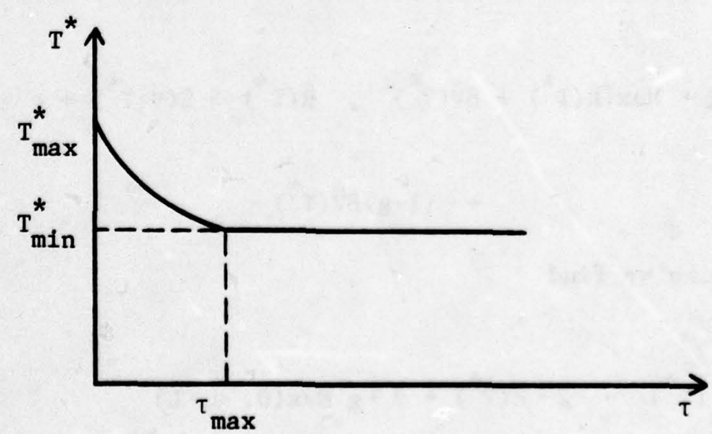

Figure 6.8 . T^* as a function of τ .

information anticipating more free information from a coming decision.

IV. Prompt but imperfect information from decisions.

Finally we study the case where prompt but imperfect information is learned from decisions. This is the case where the outcome of a decision is revealed promptly but the state cannot be perfectly observed from the outcome. Let us assume that each time a decision is made, the level of our information jumps to $\theta \cdot v_o$, where v_o is the level of information immediately after a perfect observation and θ is a constant less than 1 (Fig. 6.9). We assume that the information from the decision perishes

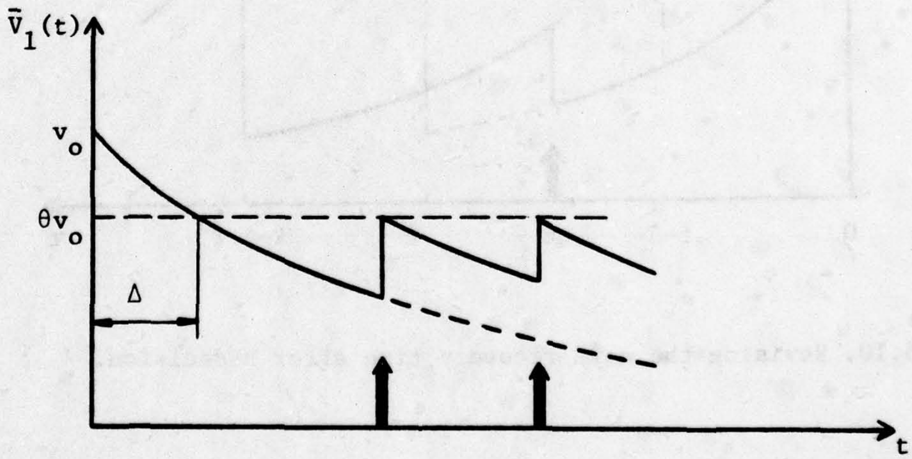

Figure 6.9. Prompt but imperfect information from decisions.

exactly like the information from an observation. Let Δ denote the interval after each observation in which the value of information is greater than θv_o (Fig. 6.9). If a decision happens in this interval, we assume that the information from this decision will not improve our information and we will remain in the same information curve. As a result, the planned time for the next recovery will not change after a decision in this interval. If a decision occurs at $t > \Delta$ (Fig. 6.10), the information from this decision (after it is revealed) is tantamount to an observation

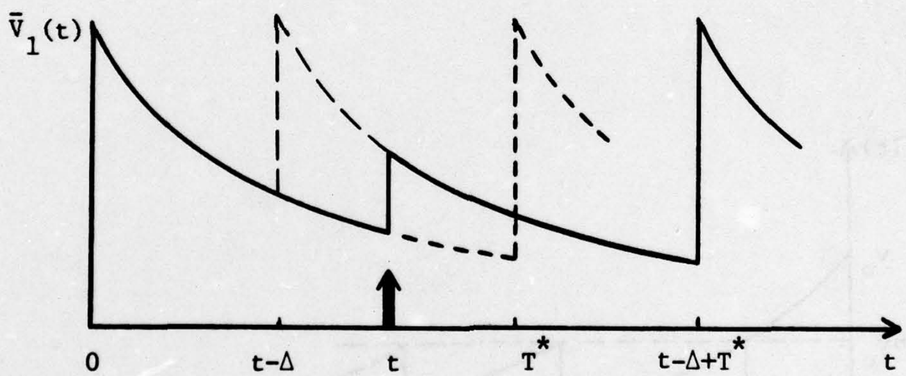

Figure 6.10. Revising the next recovery time after a decision.

at $t - \Delta$. Therefore, the next recovery time must be revised to $t - \Delta + T^*$ (Fig. 6.10)

The optimality condition for T^* can be obtained, as before, by setting the expected payoffs (at T^*) of buying information and waiting

(at T^*) equal. The decision tree is the same as in the previous cases, except for the payoff when we wait at T^* and the decision happens at T^* (Fig. 6.11). In this case we will not buy information at $T^* + 1$, but will rather wait until $T^* - \Delta + T^*$, and the expected payoff is $\bar{V}'(T^*)$.

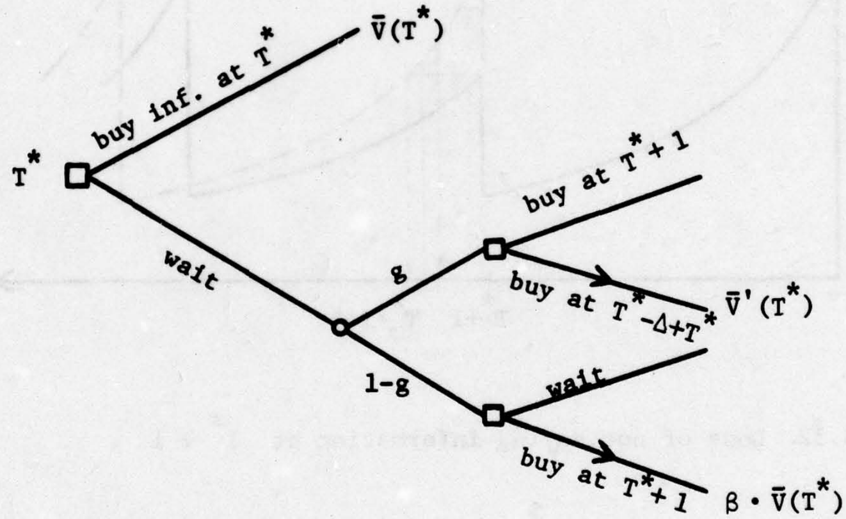

Figure 6.11. Expected payoffs at T^* with prompt but imperfect information from decisions.

$\bar{V}'(T^*)$ can be found from Fig. 6.12. Notice that the payoff from time $T^* + 1$ on (with the information from the decision) is the same (but starting

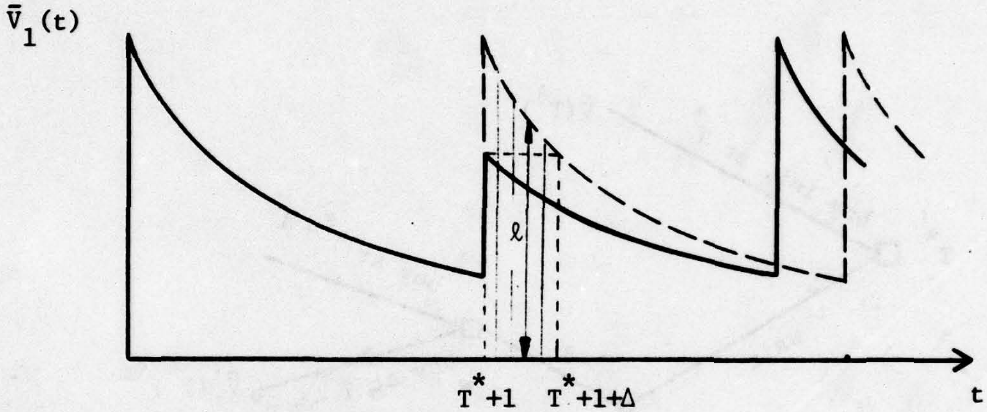

Figure 6.12. Loss of not buying information at $T^* + 1$.

Δ units earlier) as the expected payoff from time $T^*+1+\Delta$ on with an observation at T^*+1 . Therefore we can write

$$\bar{V}'(T^*) = \bar{R}(T^*) + \beta^{-\Delta-1}[\bar{V}(T^*) + c - L] \tag{6.6}$$

where L is the present value (at T^*+1) of the equivalent loss (ℓ) , as shown in Fig. 6.12. The optimality condition is:

$$V(T^*) = g \cdot \bar{V}'(T^*) + (1-g)\beta\bar{V}(T^*) \tag{6.7}$$

Substituting for $\bar{V}'(T^*)$ in (6.7) from (6.6) and after simplifying we find

$$\bar{V}_u(T^*)\left[1-g(1+\beta^{-1}+\beta^{-2}+\ldots+\beta^{-\Delta-1})\right] = g\,\bar{R}(T^*) + g\beta^{-\Delta-1}(c-L) \tag{6.8}$$

This condition is similar to the previous optimality conditions, but it is somewhat complicated. The benefits of information from the decisions are reflected on both sides of the equation, and therefore cannot be interpreted as easily as before.

T^* as a function of θ has the shape of Fig. 6.13. For θ less than some θ_o, the recovery period is constant (T^*_{min}). This corresponds

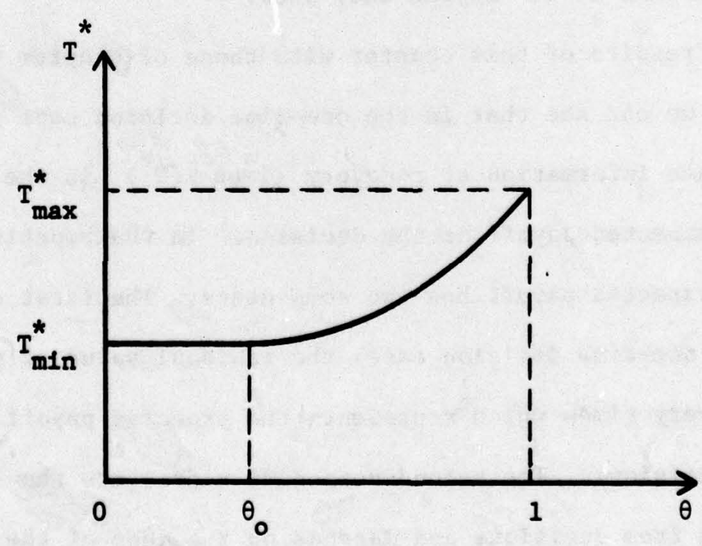

Figure 6.13. T^* as a function of θ.

to the case where the information from the decision is less than the residual value of information at the recovery time (with no information from decisions). Therefore, there is no value to this information and T^*_{min} is the same as T^* with no information (case I). As θ increases above θ_o, the information from the decisions is useful and T^* increases (we buy less information anticipating free information from a coming decision). For $\theta = 1$, T^* reaches its maximum (T^*_{max}) which is the same as T^* in case II. Recall that the same results were obtained in case III, where we found T^* as a function of the delay τ. On the other hand, we argued in case II that T^* for that case ($= T^*_{max}$) is the same as T^* for the one-time decision case. From these observations we can state that:

<u>Proposition 6.2</u>. If the decision is repeated in time we always buy more information than if it happens only once.

Comparing the results of this chapter with those of Chapter 4 (one-time decision case) we can see that in the one-time decision case the residual value of past information at recovery times $\bar{R}(T^*)$ is the measure of the net expected payoff of the decision. In the repetitive case, however, the expected payoff has two components. The first component is (like the one-time decision case) the residual value of past information at recovery times which represents the expected payoff without learning from decisions. The second component represents the benefits of learning from decisions and depends on the cost of the information and the probability of the occurrence of the decision. In the one-time decision case an increase in T^* implies less payoff (because $\bar{R}(T^*)$ decreases). This may not be true for the repetitive

case because an increase in T^* may have resulted from learning more information from decisions and, therefore, may imply a greater payoff.

6.3 Summary

In this chapter we have extended the results obtained for a one-time contingent decision (Chapter 4) to the case, where the contingent decision may be repeated in the future. We have noted two important differences from the one-time decision case: (1) In the repetitive-decision case each piece of information may be used for more than one decision and, therefore, information is potentially more valuable; and (2) there may be an opportunity to learn from each decision about the state at the time of the decision, and use this free information for future decisions. Four cases have been studied corresponding to four types of information learned from decisions: no information from decisions, perfect information from decisions, perfect but delayed information from decisions, and prompt but imperfect information from decisions.

The optimality condition is found for each of the above cases. It was noted that if this condition is written in terms of the uniform-payment equivalent of the future payoffs, the optimality condition will be very similar to the condition for the one-time decision case. In all cases we found that for the optimal recovery policy the residual value of the past information immediately before a recovery is a measure of the expected payoff of each decision (as it was in the case of a one-time decision). However, there is an additional component to the payoff, which is due to the free information learned from decisions. This component appears in the optimality condition as a separate term,

and gives us the exact benefits of learning from decisions (provided that information is recovered optimally).

It was shown that in the repetitive-decision case we always buy more information than in the one-time decision case, because information may be used for more than one decision and is, therefore, more valuable. However, the more information we learn form the decisions, the less information we will buy (anticipating free information from a forthcoming decision). When perfect information is learned from each decision, we will buy information at the same rate as in the one-time decision case.

An interesting question regards the manner in which the time of the next recovery must be revised after a decision occurs and free information is learned (or will be learned) from the outcome of the decision. It was noted that the next recovery time must, generally, be changed to a time, when the information from the decision has reduced in value to the level just before a recovery. There are some exceptions to this rule, however. For example, when we learn perfect but delayed information from decisions, we have to wait before the information from the decision is revealed. During this waiting period we may be low on information and have to bear a loss, if the decision occurs. If this loss is larger than the expected benefit of the information from the decision, we would be better off ignoring the information from the decision and buy information at the previously planned time.

CHAPTER 7

SUMMARY

7.1 Conclusions

In the first part of this research the process of information outdating was analyzed in time. We determined how the process depends on the characteristics of the dynamic environment as well as on the decision for which the information is used, and on the properties of the information itself.

Assuming a quadratic payoff function the value of a given piece of information was found as a function of its age. The result was then put in a form which shows separately the effect of various factors on the value of information. We observed that the dynamics of the state (environment) are the main determinants of the dynamics of information and manifest themselves through changes over time of the mean of the state, given the state at the observation time. If the state has a normal distribution, this is the same as the changes over time of the correlation coefficient of the current state with the state at the time of the observation.

The information structure influences the value of the information through the variance of the posterior mean of the state, given fresh information. The payoff function is represented by a single matrix which is determined from the coefficient matrices of the payoff function. Although the information structure and the payoff function are represented by constant matrices, they can drastically influence the dynamics of the information.

The equation for the value of the information indicates that the outdating of information is a complicated process with a variety of patterns. Specific results were obtained for linear Markovian systems. It was shown that the eigenvalues of the state-space matrix of the system play a major role in determining the pattern of the information outdating process. We found that the value of the information may increase or may even oscillate with time. This is a counter-intuitive result, especially for a Markovian system. It was shown, however, that if the information is perfect its value always decreases with time, regardless of the dynamics of the state or the parameters of the payoff function. Other cases of information perishing were also identified. For these cases, bounds for the rate of information perishing were determined by the smallest and the largest eigenvalues of the state-space matrix of the system.

The information outdating process for autoregressive systems was also investigated. The results gave insight into the process of information outdating, in particular in cases where the value of the information is enhanced or oscillates with time.

In the second part of this research we have investigated the optimal policies for recovery (updating) of information when anticipating decisions at uncertain times. Both the a priori and the a posteriori optimal information recovery policies were studied. The case of a one-time contingent decision was initially analyzed. The results were then extended to the case of repetitive decisions in time.

A necessary condition for an optimal information recovery policy was found for the general case with the following interpretation: at an optimal information recovery time, the expected marginal loss of buying information slightly later is equal to its expected marginal gain in the future. For a one-time decision case with an infinite horizon, the a priori optimality condition reduces to a simple form and states that the net expected payoff of the decision is equal to the expected payoff, if the decision were to occur when the information has its lowest value (immediately before each recovery). The extra payoff of the decision, were it to occur at other points in time, will be just enough to pay for the cost of the information.

We have shown that the information recovery problem can be formulated using, as our basic variable, either the "residual value of past information," or the "value of new information" at each point of time. The two variables are intimately related and may be regarded as duals. It was found, however, that the latter is more appropriate, especially for studying the a posteriori information recovery policies.

In addition to simplifying the optimality condition, the new variable permitted an important observation about the optimal information recovery policies, concerning the following question: Under what conditions is the optimal information recovery time independent of the result of the previous information recoveries? Under such conditions, the a posteriori and the a priori information recovery policies are identical and, therefore, the computation of the optimal policy is

greatly simplified. Using the new variable we showed that if the (total) value of new information, given that it is bought only once, is always independent of the result of the previous observations, then each optimal recovery time is independent of the result of the previous observations. This result is general and requires no condition on the state, the decision, or the information structure. One important example of the above property occurs when the payoff function is quadratic and the state has a normal distribution.

In the repetitive-decision case there are two distinct differences from the one-time decision case: (1) each piece of information may be used for more than one decision and, therefore, the information is potentially more valuable; and (2) there may be an opportunity to learn about the state from each decision and use this free information for future decisions. Several cases corresponding to various types of information learned from the decisions were studied. We found that for an optimal information recovery policy the residual value of the information immediately before a recovery is a measure of the expected payoff of each decision (as it was in the case of a one-time decision). However, in view of the free information learned from each decision the payoff contains an additional component. This component appears in the optimality condition as a separate term and yields the exact benefits of learning from the decisions.

It was shown that in the repetitive-decision case we always buy more information than in the one-time decision case (because information

may be used for more than one decision and is, therefore, more valuable). However, the more information learned from the decisions, the less information we will buy. When perfect information is learned from each decision, information will be bought at the same rate as in the one-time decision case.

Finally, the effect of risk aversion on the optimal information recovery policies was studied briefly. The optimality condition for an expected-value decision maker was extended to the case of a risk-averse decision maker. It was found very difficult to make any comparison between the two cases, however, because the payoff lotteries are different for the two decision makers. Nevertheless, we found evidences which suggest that a risk-averse decision maker would buy less information than an expected-value decision maker, when there is uncertainty regarding the time of the decision. Whether (or under what circumstances) this is in fact true, remains uncertain.

7.2 Suggestions for Future Research

This research has been a first step in the investigation of information in a dynamic framework. Our emphasis has been on the formulation of the problem and the development of simple theoretical models which give insight into the problem, and provide a basis for future research. There is thus considerable potential for future research to develop more complete models based on real-world examples, and to establish more concrete criteria for rationalizing the information production process. Opportunities for future research exists in several areas, a few of which are suggested below.

One immediate extension of the work in this study will be to consider a changing payoff function. We have assumed, for simplicity, that the payoff function does not vary with time. This may not be the case in many real-world situations. When the payoff function changes with time, its effects on the information perishing process, and also on the optimal information recovery policies are unknown.

A second area would be to extend the results to a nonstationary state (environment) which is important since this is typically the case in practice. There is no fundamental difficulty in extending the results to nonstationary systems, although this will increase the burden of computation.

An important extension of this study concerns the model of information acquisition. In our model we have assumed that all the uncertain states are simultaneously observed. A more realistic model would permit the observation of various uncertain states at different points in time, which implies that some states could be observed more frequently than others.

Another extension would allow the underlying model of the dynamic environment to be updated. In this study, the model of the environment is considered to be exogenous. The model can be endogenously determined from previously accumulated information, and revised as new information is acquired. In other words, the model itself is also updated. This is clearly a more appropriate model for an applied study.

APPENDIX

NOTATION

$s(t)$	one-dimensional state variable at time t .
d	one-dimensional decision variable.
$\underline{s}(t)$	vector of state variables at time t .
$\underline{s}'(t)$	transpose of $\underline{s}(t)$.
$\underline{d}$	vector of decision variables .
$v_1(\underline{s}(t),\underline{d})$	payoff of the decision $\underline{d}$ (at time t), given $s(t)$.
$\mathscr{E}$	prior knowledge
$\{\underline{s}(t)\|\mathscr{E}\}$	joint probability distribution over $\underline{s}(t)$, given $\mathscr{E}$.
$\langle \underline{s}(t)\|\mathscr{E}\rangle$	vector of expected value of $\underline{s}(t)$, given $\mathscr{E}$.
${}^{V}\langle \underline{s}(t)\|\mathscr{E}\rangle$	covariance matrix of $\underline{s}(t)$, given $\mathscr{E}$.
$\eta(\underline{s}(t))$	information structure (observation) on $s(t)$.
$\underline{z}(t)$	vector of observation of $\underline{s}(t)$ $\big(\underline{z}(t) = \eta(\underline{s}(t))\big)$
$\underline{y}(t)$	information (data) available at time t .
$\{\underline{s}(t)\|\underline{y}(t),\mathscr{E}\}$	joint probability distribution over $\underline{s}(t)$, given $\underline{y}(t)$ and $\mathscr{E}$.
$\langle \underline{s}(t)\|\underline{y}(t),\mathscr{E}\rangle$	vector of expected value of $\underline{s}(t)$, given $y(t)$ and $\mathscr{E}$.
${}^{V}\langle \underline{s}(t)\|\underline{y}(t),\mathscr{E}\rangle$	covariance matrix of $\underline{s}(t)$, given $\underline{y}(t)$ and $\mathscr{E}$.
$V_{\eta(t-\tau)}(t)$	expected value at time t of $\eta(\underline{s}(t-\tau)$.
$V_{\eta(\tau)}$	$V_{\eta(t-\tau)}(t)$ for a stationary state variable.
$V'_{z(t_o)}(t)$	expected value at time t of $z(t_o)$.

$V'_{\underline{z}_o}$	$V'_{\underline{z}(0)}(t)$ for a stationary state .
$\rho(\tau)$	rate of information perishing.
$\underline{\tilde{s}}(t)$	posterior mean of $\underline{s}(t)$, given $\underline{y}(t)$.
$\Sigma_{\underline{\tilde{s}}_o}$	covariance matrix of $\underline{\tilde{s}}(o)$.
$R(\tau)$	coefficient of linear approximation of $s(t)$ as a function of $s(t-\tau)$.
Z	set of observations $\underline{z}$.
ZfZ'	set Z of observation $\underline{z}$ is "finer" than set Z' of observation $\underline{z}'$.
ZgZ'	set Z of observation $\underline{z}$ is "garbled" into set Z' of observation $\underline{z}'$.
$\eta \succcurlyeq \eta'$	η is "more informative" than η' .
E	precipitating event in a contingent decision.
c	cost of each observation.
g_t	probability of the decision occurring at time t .
Z_t	result of the previous observations at time t .
$V_1(t,Z_t)$	expected payoff of the decision at time t , given Z_t .
$\bar{V}_1(t)$	a priori expected payoff of the decision at time t .
$\bar{R}(t)$	residual value of information at time t .
$V_t(t_1,Z_o)$	maximum expected payoff from time t on, with the next recovery at t_1 , given Z_o .
$\bar{V}_t(t_1)$	a priori expected payoff from time t on, with the next recovery at t_1 .
$V(t,Z_o)$	maximum expected payoff from time t on, with recovery at t , given $Z_o (= V_t(t,Z_o))$.

$\bar{V}(t)$	a priori maximum expected payoff from time t on, with recovery at t $(= \bar{V}_t(t))$.
$\bar{V}(T)$	a priori net expected future payoff, with information recovery period T .
$V_1'(t,Z_o)$	expected payoff of the decision at time t , with recovery at t , given Z_o .
$\bar{V}'_1(t)$	a priori expected payoff of the decision at time t , with recovery at t .
T^*	a priori optimum information recovery period.
$t^*(Z_o)$	optimum time for the next information recovery, given the result of the previous observations, Z_o .
$\bar{U}(L)$	expected utility of the lottery L .
$W_1(t,Z_o\|T)$	value of new information (obtained at T) for making the decision at t , given Z_o .
$W(t,Z_o)$	net value of all the future purchases of information with the first one at t , given Z_o .
$W^1(T,Z_o)$	total net value of information bought at T , with no more information in future, given Z_o .
$\bar{V}_u(T)$	uniform-payment equivalent of $\bar{V}(T)$.
β	rate of discount .

REFERENCES

[1] Agnew, Carson E., Proximal Decision Analysis with Imperfect Information, Department of Engineering-Economic Systems, Stanford University, May 1975.

[2] Anderson, Theodore W., An Introduction to Multivariate Statistical Analysis, New York, John Wiley and Sons, 1958.

[3] Bellman, Richard, Introduction to Matrix Analysis, Second Edition, McGraw-Hill Book Company, 1970.

[4] Box, George E., and Gwilym M. Jenkins, "Time Series Analysis, Forecasting, and Control,' San Francisco, Holden-Day, 1976.

[5] Grum, Allen F., Information Perishing and Replenishment, Ph.D. Dissertation, Department of Engineering-Economic Systems, Stanford University, 1976.

[6] Howard, Ronald A., Information Value Theory, IEEE Transactions on Systems Science and Cybernetics, Vol. SSC-2, No. 1, August 1966, pp. 22-26.

[7] Howard, Ronald A., Dynamic Probabilistic Systems, Vol. 2., John Wiley & Sons, Inc., 1971.

[8] Howard, Ronald A., and Matheson, James E. (eds.), Readings in Decision Analysis, Stanford Research Institute, Menlo Park, 1976.

[9] Marschak, Jacob, Information Economics Reconsidered, Working Paper No. 149, Western Management Science Institute, University of California, Los Angeles, June 1969.

[10] Marschak, Jacob, Economics of Information Systems, Journal of the American Statistical Association, Vol. 66, No. 333, March 1971, pp. 192-219.

[11] Marschak, Jacob, and Roy Radner, Economic Theory of Teams, New Haven, Yale University Press, 1972.

[12] Marschak, Jacob, and Koichi Miyosawa, Economic Comparability of Information Systems, International Economic Review, Vol. 9, No. 2, June 1968, pp. 137-176.

[13] McDonough, Ardin M., Information Economics and Management Systems, McGraw Hill Book, Co., Inc., New York, 1963.

[14] McGuire, C. B., Comparisons of Information Structures, in C. B. McGuire, and Roy Radner (eds.), Decision and Organization, North-Holland Publishing Company, 1972.

[15] Merkhofer, Miley W., The Value of Information Given Decision Flexibility, Ph.D. Dissertation, Department of Engineering-Economic Systems, Stanford University, June 1975.

[16] Raiffa, Howard, and Robert Schlaifer, Applied Statistical Decision Theory, The M.I.T. Press, Cambridge, Massachusetts, 1977.